arXiv:1407.5298v2 [cs.DS] 26 Nov 2015

How the Experts Algorithm Can Help Solve LPs Online*

Anupam Gupta† Marco Molinaro‡

November 30, 2015

Abstract

We consider the problem of solving packing/covering LPs online, when the columns of the constraint matrix are presented in random order. This problem has received much attention and the main focus is to figure out how large the right-hand sides of the LPs have to be (compared to the entries on the left-hand side of the constraints) to allow $(1+\varepsilon)$-approximations online. It is known that the right-hand sides have to be $\Omega(\varepsilon^{-2}\log m)$ times the left-hand sides, where m is the number of constraints.

In this paper we give a primal-dual algorithm that achieve this bound for mixed packing/covering LPs. Our algorithms construct dual solutions using a regret-minimizing online learning algorithm in a black-box fashion, and use them to construct primal solutions. The adversarial guarantee that holds for the constructed duals helps us to take care of most of the correlations that arise in the algorithm; the remaining correlations are handled via martingale concentration and maximal inequalities. These ideas lead to conceptually simple and modular algorithms, which we hope will be useful in other contexts.

1 Introduction

Sequential decision-making is one of the most important topics in Operations Research. While several sequential decision models are available, the common feature is that items/requests come over time and at a given time only partial information about the future items/requests is known. Decisions need to be made over time only with the information available at that point. Moreover, typically decisions are irrevocable.

While such problems have been extensively studied in the worst-case competitive analysis [BN09, ABFP13], in order to avoid the pessimistic achievable guarantees a lot of recent research has focused on the *random-order* model, which can be broadly described as follows: while the underlying instance is arbitrary (i.e. chosen adversarially), the items/requests come in a **random order** (i.e. the adversary cannot control the order in which the instance is presented). Several sequential decision problems have been considered in this model, starting from a classic maximization version of the secretary problem [Dyn63], to single-knapsack problems [Kle05, BIKK07], the AdWords problem [MSVV07, GM08, DH09], matroid secretary problems [BIK07, BIKK08], and packing LPs [FHK+10, AWY14, MR12, KRTV13, AD15, ESF14].

In this paper we consider a broad class of such problems, namely that of solving online packing/covering multiple-choice LPs; this class includes most of the examples above.

*A preliminary version of this work appeared in ESA 2014.

†Computer Science Department, Carnegie Mellon University, Pittsburgh, PA. Research partly supported by NSF awards CCF-1016799 and CCF-1319811, and by a grant from the CMU-Microsoft Center for Computational Thinking.

‡EWI, TU Delft.

1.1 Model and prior results

The offline version of the *packing/covering multiple-choice LPs* we consider have the form

$$\begin{aligned} \max \ & \textstyle\sum_{t=1}^{n} \pi^t x^t \\ st \ & \textstyle\sum_{t=1}^{n} A^t x^t \le b \\ & \textstyle\sum_{t=1}^{n} C^t x^t \ge d \\ & x^t \in \blacktriangle^k \quad \forall t \in [n] \end{aligned} \tag{PCMC}$$

where all the data π^t, A^t, C^t, b, d is **non-negative** and $\blacktriangle^k$ denotes the "full simplex" $\{x \in [0,1]^k : \sum_j x_j \le 1\}$. Note that the variables $\{x_1^t, x_2^t, \ldots, x_k^t\}$ in each block t are bound together by $\sum_j x_j^t \le 1$ — this is the "multiple-choice" setting, which is useful in many contexts. In the case $k = 1$, the multiple-choice model degenerates to the standard online packing-covering LP model where a single variable x^t is revealed at each time, and must be assigned a fractional value in $[0, 1]$.

In the *random-order model*, the number n of columns and the right-hand sides b and d are known a priori but the columns of the LP are revealed one by one in random order, that is, their arrival order is chosen uniformly at random out of all possible permutations. When a column-block (π^t, A^t, C^t) is revealed, the corresponding x^t variables need to be irrevocably set before the next column-blocks are revealed. The goal is to obtain a feasible solution for the LP that tries to obtain as much value (in expectation or with high probability) as the optimal **offline** solution. This is exactly the model we consider here.

Packing/covering multiple-choice LPs form a broad class of problems. Special cases of this model with **packing-only** constraints (i.e. without the constraints $\sum_{t=1}^{n} C^t x^t \ge d$) have been extensively studied [Dyn63, Kle05, BIKK07, MSVV07, GM08, DH09, FHK$^+$10, AWY14, MR12, KRTV13, AD15, ESF14]. These models have a vast range of applications, like online advertisement, online routing, and airline revenue management.

Most of these results show that one can obtain online an $(1-\varepsilon)$-approximation for the LP as long as the right-hand side of the constraints are sufficiently (depending on ε) larger than any entry in the left-hand side matrix. This direction started with the work of Kleinberg [Kle05], who considered the case without multiple choices (i.e. $k = 1$) and with a single cardinality constraint $\sum_t x^t \le b$. He showed a $(1-\varepsilon)$-approximation in expectation as long as $\frac{b}{\max_t a^t} \ge \Omega(\epsilon^{-2})$, which is best possible. Subsequently, Devanur and Hayes considered the AdWords problem [DH09], which was then extended to general packing multiple-choice LPs [FHK$^+$10, AWY14, MR12]. These results provide an $(1-\varepsilon)$-approximation as long as the right-hand sides/item sizes ratio satisfies $\frac{\min_i b_i}{\max_{t,i,j} A_{ij}^t} \ge \Omega(\varepsilon^{-2}) \min\{m^2 \log(m/\varepsilon), m \log(n/\varepsilon)\}$, where m is the number of packing constraints and n the number of column-blocks. Agrawal et al. [AWY14] also showed that $\Omega(\varepsilon^{-2} \log m)$ ratio is required for $(1-\varepsilon)$-approximations, which was believed to be the right answer.

However, despite the increasing interest of solving LPs online, to the best of our knowledge no guarantee is known for online optimization of LPs with **covering** constraints, even in the weaker i.i.d. model (see [MR12] for discussion of random order vs. i.i.d.)[1]

1.2 Our results

Packing/covering multiple-choice LPs. Our main contribution is the first set of algorithms with provable guarantees for packing/covering multiple-choice LPs in the random-order model. Our algorithm handles general packing/covering multi-choice LPs and achieves the optimal guarantee even

[1] [DJSW11] presents an online for the **feasibility** version of packing-covering multi-choice LPs in the i.i.d. model.

when only packing constraints are present. For example, in case an estimate of the optimal value is present we have the following guarantees.

Informal Theorem 1.1 (Packing/covering LPs)**.** *Consider a feasible packing/covering multiple-choice LP in the random-order model. Given ε, δ, assume that each right-hand side in the LP is at least $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ times the maximum coefficient of the left-hand side in its row, and that the optimal value is at least $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ times each individual item's profit. Then given an (under)estimate O^* for the optimal value of the offline LP, the algorithm **LPviaLB** computes an ε-feasible solution online that achieves a value of $(1-\varepsilon)\,O^*$ with probability $1-\delta$.*

(For the formal version, see Theorem 4.1.) Moreover, the algorithm is very efficient and does not need to solve any (offline) Linear Programs. Also, as is common in this context, our algorithm produces *integer* solutions that compare favorably to the optimal offline fractional solution.

The assumption of the availability of an estimate for the optimal value is often a reasonable one (e.g. when enough historic knowledge of the problem directly provides such an estimate). Nonetheless, we show how to construct such an estimate as the columns arrive online. However, to obtain good estimates we need to consider *stable* LPs, which informally means that allowing solutions to violate the covering constraints by a small factor does not increase the optimal value by much (see Definition 5.2, and also Lemma 5.3 for a simple sufficient condition for stability). We note that a packing-only LP is vacuously stable.

Informal Theorem 1.2 (Packing/covering LPs II)**.** *Consider a feasible, stable, packing/covering multiple-choice LP in the random-order setting. Given ε, δ, assume that each right-hand side in the LP is at least $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ times the maximum coefficient of the left-hand side in its row, and that the optimal value is at least $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ times each individual item's profit. Then with probability at least $1-\delta$ the algorithm **DLA** computes an ε-feasible solution online that has value at least $(1-\varepsilon)$ times the optimal value of the offline LP.*

(The formal version appears as Theorem 6.2.) The high-level idea to obtain an estimate of the optimum is reasonably standard at this point ([Kle05, DJSW11, AWY14, MR12, ESF14]: one solves an LP with the columns seen so far to obtain an estimate, which is updated at doubling times. However, the presence of covering constraints introduces additional technical difficulties, and we need the assumption of stability of the LP to control the variance of estimate produced (this issue is further discussed in §1.4.1).

Packing LPs. Informal Theorem 1.2 assumes that individual item profits are small compared to the optimal value. While this is often reasonable, can we remove this assumption? Our final result removes all assumptions about the magnitude of the profit whenever we are dealing with packing-only LPs (and without multiple-choice constraints of the form $\sum_j x_j^t \leq 1$). Obtaining good estimates of opt without any assumptions on the item values in our setting requires significantly new ideas (also discussed further below in §1.4.1). (The formal version of the following theorem appears as Theorem 7.11.)

Informal Theorem 1.3 (Ultimate packing LP)**.** *Consider a packing LP in the random-order setting. Given $\varepsilon > 0$, suppose the right-hand sides are $\Omega(\frac{\log(m/\varepsilon)}{\varepsilon^2})$ times any left-hand side entry. Then the algorithm **mDLA** computes a solution online with expected value at least $(1-\varepsilon)$ times the optimal value of the offline LP.*

Generalized load balancing. The main building block for our algorithms for packing/covering LPs is the ability to solve the following *generalized load-balancing problem*: We have a set of m

machines. For each arriving job t, there are k different options on how to serve it, and the j^{th} choice requires some amount of processing on each of the machines, given by the vector $(a^t_{1j}, a^t_{2j}, \ldots, a^t_{mj})$. When the job arrives, we must (fractionally) choose one of these options. The goal is to minimize the makespan, i.e., the maximum total processing assigned to any machine. Notice that when the matrices A^t's are diagonal $m \times m$, the j^{th} choice places a^t_{jj} load only on the j^{th} machine, and hence captures the classic problem of scheduling on unrelated machines.

Informal Theorem 1.4 (Load-balancing)**.** *Consider the load-balancing problem where the jobs arrive in random order. Given $\varepsilon > 0$, if the optimal makespan λ^* is at least $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ times the maximum amount of processing required by any option for any job, the algorithm **expertLB** constructs an online solution with makespan at most $(1+\varepsilon)\lambda^*$ with probability $1-\delta$.*

(The formal version in Theorem 3.3 proves this with respect to a generic low-regret online learning algorithm for the prediction-with-experts problem. For readability, here we instantiate it with parameters from the Multiplicative Weights Update algorithm.) The case where the jobs are **i.i.d.** draws from some unknown distribution was considered in [DJSW11, Section 3], where with probability $1-\delta$, a $(1+\varepsilon)$-approximation to the optimum makespan is obtained assuming the optimum is $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ times the maximum processing a^t_{ij}. Our theorem above shows that the same guarantee holds in the more general random-order model, answering an open question from [DJSW11].

In fact, we can handle processing requirements that can be both positive or negative, as long as for each machine the jobs are mostly positive or mostly negative (see Definition 3.1). Handling negative processing requirements is crucial for addressing covering constraints when we **reduce** packing/covering LPs to the generalized load-balancing problem. We note that [DJSW11] also considered this generalized load-balancing problem in connection with packing/covering LPs; however, they do not consider negative processing times and thus do not provide a direct reduction between the problems.

1.3 The Road-map

Our algorithmic development is modular: before elaborating on our technique, let us lay out the plan of attack.

- In §3 we show our results for generalized load-balancing on random-order arrival streams via online learning algorithms.
- In §4 we give a simple reduction that reduces solving packing/covering LPs *with an estimate of the optimal value* to generalized load-balancing, and hence prove Theorem 1.1.
- In §5 we show how to estimate the optimal value for packing/covering LPs.
- In §6, we combine ideas from §4 and §5 to prove Theorem 1.2, our main result about packing/covering LPs.
- Finally, in §7 we refine the results of §6 for packing-only LPs: we remove the requirement that the individual item profits should be small, and get Theorem 1.3.

1.4 Our Techniques

Decoupling primal and dual via online learning. Dual-based method are very important in optimization under uncertainty and have been used extensively in the Operations Management literature [Sim89, Wil92, TR98]. One advantage of primal-dual methods is that they typically have fast update times, and avoid solving an LP at every step of the process. In addition, the computed duals can be used as a good heuristic to estimate the importance of each constraint and to perform sensitivity analysis.

Most of the works on packing LPs in the random-order model [DH09, FHK$^+$10, AWY14, MR12] use the following form of the primal-dual approach (the primal-only algorithm of [KRTV13] being a

notable exception)[2]: a good dual LP solution is estimated based on the columns seen so far, then the reduced costs defined by this dual are used to decide how to set the primal variables. While being a natural idea, analyzing it poses the following crucial difficulty: the primal and dual solutions, by construction, are correlated in a non-trivial manner. In order to deal with these correlations, all of the above analyses resort to a massive union bound, which is the root of the extra $\log n$ and m factors over the right ratio bound $\Omega(\epsilon^{-2} \log m)$.

One of our contributions is an alternative algorithmic construction where the primal and the dual are only loosely coupled. For that, we construct duals using a generic *regret-minimizing online learning* algorithm (which offers *adversarial* guarantees); in the primal, we use a greedy strategy which ensures that, with respect to these duals, we are doing at least as well as the optimal offline solution.

The fact that the guarantees for the dual hold in the adversarial setting, together with the greedy primal step, reduces the analysis of the algorithm to that of comparing the dual solution constructed to the *offline optimal* solution (see §3.4). Since the latter is a *fixed* solution (modulo the permutation of the columns), it is only loosely correlated with our constructed dual; in fact, correlations only arise because in the random-order model we are sampling *without* replacement (in particular, in the weaker i.i.d. model there are no such correlations, see §3 up to equation (3.1)).

As far as we know, this is the first explicit black-box connection between online learning and optimization in the random-order model (see §1.5 for independent and subsequent work); in the offline packing/covering LP setting such a connection was implicit in, e.g., [PST95, You01, GK07] and made explicit in, e.g., [Kha04, Haz06, AHK12, CHW12]. In the simpler i.i.d. setting, [DJSW11] uses multiplicative-weights based dual updates. This is reminiscent of the Multiplicative Weight Update algorithm for online learning. Our analysis abstracts out the reliance on the details of that technique, giving a clean conceptual explanation of what is happening. It also allows us to use *any* experts online learning algorithm in our reductions, where the quality of the LP solution depends on the parameters in the rate of regret. Importantly, this isolation between the analysis of the online learning algorithm and the comparison of our dual solution and the offline optimum is instrumental in dealing with the dependencies arising in the random-order model.

Handling dependencies in random permutations. Despite the loose coupling of primal and dual solutions, the correlations arising from sampling *without* replacement cannot be immediately handled: the known bounds on the statistical distances between sampling with and without replacement (e.g. [DF80]) do not seem to be strong enough, and a simple approach requires taking a union bound across all time steps, leading again to extra $\log n$ factors. To overcome this, we instead use a *maximal inequality* for sampling without replacement to compare sampling with/without replacement at *every* time step while avoiding a union bound, together with martingale concentration inequalities. The arguments we develop are quite general and may find use in other problems in the random-order model.

1.4.1 Estimating OPT: Challenges and Ideas

In our approach we need to estimate (with high probability) the optimal value opt of the offline LP. *Under the assumption that the individual item profits are small compared to* opt, such estimate can be obtained by solving the LP comprising the first few columns seen in random order, and scaling its optimal value appropriately (actually for packing/covering LPs even this case is not that easy, and we require stability from the LP to provide guarantees for such estimates). However, without assuming anything on the individual item sizes, as we do for the packing-only case in §7, the high-valued items make this estimate have too high a variance and precludes the high-probability guarantees required.

[2]Primal-dual methods of a slightly different flavor are also extensively used in the worst-case model of online algorithms [BN09].

In fact **all** previous approaches that use explicit estimates of opt make the additional assumption that each item profit is at most $O(\frac{\varepsilon^2}{\log m})\,\mathsf{opt}$ [DJSW11, ESF14, AD15][3] To remove such assumptions, we show that one can simply pick all items with sufficiently high value and add them to an $(1-\varepsilon)$-approximation for the remaining LP (which then does not contain high-valued items) and this only incur a small loss in total profit (see §7, in particular Lemma 7.3). Unfortunately we also need to estimate the *threshold* of what constitutes high enough value; this is however an easier task, since we only need weak guarantees for this estimate. This estimate is obtained by again looking at the LP comprising the first few columns seen in random order but *removing the K highest valued items* in this sample; this provides sufficient control on the remaining item values for the analysis to go through.

1.5 Independent and subsequent work

Independent to our work, Kesselheim et al. [KRTV13] present a primal algorithm that achieves the first optimal guarantee for **packing-only** multiple-choice LPs. More precisely, if d is the *column-sparsity* of the LP (i.e., maximum number of non-zero entries in each column), they obtain an $(1-\varepsilon)$-approximation as long as the right-hand sides/item sizes ratios are at least $\Omega(\epsilon^{-2}\log d)$. The dependence on the column-sparsity d rather than the number of constraints m (and the fact they handle multi-choice constraints) makes their result quantitatively stronger than our Theorem 1.3, while both results coincide and are optimal for $d = \Omega(m)$ [AWY14]. However, our algorithm is more efficient since it requires us to solve only a small number of offline LPs, while their algorithm solves an LP at each of the n timesteps (in the experiments presented in [ESF14], the algorithm from [KRTV13] is 2 to 3 orders of magnitude slower than algorithms that solve roughly as many LPs as ours does).

Subsequent to our work, Eghbali et al. [ESF14] provided a primal-dual algorithm for solving more general packing-type *convex* programs. They also propose a primal-dual algorithm that draws inspiration from [DJSW11]; in particular, they also use a sort of exponential update of duals. In their analysis, they also employ martingale methods and a (Doob's) maximal inequality.

Also subsequent to our work, Agrawal and Devanur [AD15] consider an even broader class of online convex programs (see the paper for proper description). Their primal-dual algorithm extends the one we present in §3 for the load-balancing problem and generalizes our connection between online (convex) learning and i.i.d./random-order convex optimization. A key element for this generalization is the use of *Fenchel duality*. One drawback is that their general result only satisfies the constraints *in expectation*. However, for the case of packing-only LPs their techniques are able to give bona fide feasible solutions also with optimal guarantees.

Finally, independently Agrawal and Devanur [AD14] used a similar technique to devise efficient an algorithm for the Bandits with concave rewards and concave knapsacks problem (Algorithm 3), again using Fenchel duality. They trace back the inspiration for this approach to [ABH11], which show how solve the Blackwell approachability problem (which at a very high-level asks for a sequence of iterates that approaches a convex set) using online convex optimization.

2 Preliminaries

For a vector $v \in \mathbb{R}^k$, define its max-norm $\|v\|_{\max} := \max_j v_j$; notice this does not involve absolute values as in the definition of the ℓ_∞ norm. Also, we use $|v| := (|v_1|, |v_2|, \ldots, |v_k|)$ to denote the component-wise absolute value of a vector. Recall that Δ^k denotes the simplex $\{p \in [0,1]^k : \sum_i p_i = 1\}$, and $\blacktriangle^k$ denotes the "full simplex" $\{x \in [0,1]^k : \sum_j x_j \le 1\}$.

[3]The requirement in [AD15] comes from the use of Lemma 14 from [DJSW11], which uses this assumption.

2.1 Prediction-with-experts problem

As mentioned before, our primal-dual algorithms will have their dual solutions constructed via an algorithm for maximization version of the classic *prediction-with-experts problem* [CBL06]. In this online problem, **adversarial** vectors $o^1, o^2, \ldots, o^n \in [-M, M]^m$ are presented one by one to the algorithm. At time t, using only the information $o^1, \ldots, o^{t-1}$ up to time $t-1$, the algorithm needs to compute a vector w^t in the simplex Δ^m; it then incurs a reward of $\langle w^t, o^t \rangle$. The goal is to maximize the total reward obtained $\sum_t \langle w^t, o^t \rangle$. (Note we have stated this problem with the domain of the adversarial vectors being the scaled one $[-M, M]^m$ instead of $[-1, 1]^m$.)

The quality of the solution constructed by the algorithm is measured using the notion of *regret*, which compares against the best *fixed* $w \in \Delta^m$ that knows all the vectors o^t in hindsight, namely $\max_{w \in \Delta} \sum_t \langle w, o^t \rangle = \| \sum_t o^t \|_{\max}$. Here we consider a slightly different parametrization of regret, which is also used in [AHK12, KV05].

Definition 2.1 (Regret)**.** *We say that an algorithm for the prediction-with-experts problem has* (α, R)-regret *if the solution* $w^1, \ldots, w^n$ *constructed satisfies*

$$\sum_{t=1}^{n} \langle w^t, o^t \rangle \geq \left\| \sum_{t=1}^{n} o^t - \alpha \sum_{t=1}^{n} |o^t| \right\|_{\max} - R.$$

This problem has been extensively studied and many algorithms achieving optimal regret are known [SS11]. The traditional $(0, M\sqrt{n \log m})$-regret is attained by several algorithms and is known to be optimal if the first parameter α in the regret is kept at 0. However, the guarantee that works best for our purposes is attained by the well-known *Multiplicative Weights Update* algorithm, which gives an (α, R)-regret with (typically) much smaller value of R at the expense of a non-zero α.

Theorem 2.2 (Theorem 2.5 of [AHK12])**.** *For every* $\epsilon \in [0, 1]$*, the Multiplicative Weights Update (**MWU**) algorithm with learning parameter* ϵ *achieves* $(\epsilon, \frac{M \log m}{\epsilon})$*-regret.*

2.2 Concentration inequalities

Another important tool for our analyses are strong concentration inequalities for non-independent processes. First we need a version of Bernstein's inequality for sampling without replacement.

Theorem 2.3 (Theorem 2.14.19 in [vdVW96])**.** *Let* $Y = \{Y_1, \ldots, Y_n\}$ *be a set of real numbers in the interval* $[0, M]$*. Let* S *be a random subset of* Y *of size* s *and let* $Y_S = \sum_{i \in S} Y_i$*. Setting* $\mu = \frac{1}{n} \sum_i Y_i$ *and* $\sigma^2 = \frac{1}{n} \sum_i (Y_i - \mu)^2$*, we have that for every* $\tau > 0$

$$\Pr(|Y_S - s\mu| \geq \tau) \leq 2 \exp\left(- \frac{\tau^2}{2s\sigma^2 + \tau M} \right)$$

Notice that we can use the fact $Y_i \in [0, M]$ to upper bound the variance σ^2 as

$$\sigma^2 \leq \frac{M}{n} \sum_i |Y_i - \mu| \leq \frac{M}{n} \left(\sum_i |Y_i| + \sum_i |\mu| \right) = 2M\mu,$$

which directly gives the following.

Corollary 2.4. *Let* $Y = \{Y_1, \ldots, Y_n\}$ *be a set of real numbers in the interval* $[0, M]$*. Let* S *be a random subset of* Y *of size* s *and let* $Y_S = \sum_{i \in S} Y_i$*. Setting* $\mu = \frac{1}{n} \sum_i Y_i$*, we have that for every* $\tau > 0$*,*

$$\Pr(|Y_S - s\mu| \geq \tau) \leq 2 \exp\left(- \frac{\tau^2}{M(4s\mu + \tau)} \right) \leq 2 \exp\left(- \min\left\{ \frac{\tau^2}{8Ms\mu}, \frac{\tau}{2M} \right\} \right).$$

We also need a *maximal* Bernstein's inequality for sampling without replacement. While we could not find such inequality explicitly stated in the literature, it easily follows from Theorem 2.3 and a Levy-type maximal inequality for exchangeable random variables of [Pru98]; its proof is presented in Appendix A.

Lemma 2.5. *Consider a set of real values $x_1, \ldots, x_n$ in $[0,1]$, and let $X_1, \ldots, X_k$ be sampled without replacement from this collection. Assume $k \leq n/2$. Let $S_i = X_1 + \ldots X_i$. Also let $\mu = \frac{1}{n}\sum_i x_i$ and $\sigma^2 = \frac{1}{n}\sum_i (x_i - \mu)^2$. Then for every $\alpha > 0$*

$$\Pr\left(\max_{i \leq k} |S_i - i\mu| \geq \alpha\right) \leq 30 \exp\left(\frac{(\alpha/24)^2}{2k\sigma^2 + (\alpha/24)}\right).$$

Finally, we need the following classic martingle concentration inequality.

Theorem 2.6 (Freedman's Inequality [Fre75])**.** *Let $X_0, X_1, \ldots, X_n$ be a real-valued martingale, with $X_0 = 0$, with respect to a filtration $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \ldots \subseteq \mathcal{F}_n$. Let $Y_t = X_t - X_{t-1}$ be the martingale differences, let M be such that $|Y_t| \leq M$ for all t and let $L = \sum_{i \leq n} \mathrm{Var}(Y_t \,|\, \mathcal{F}_{i-1})$ be the predictable variance. Then for every real numbers $\alpha, \sigma^2 \geq 0$,*

$$\Pr\big(X_n \geq \alpha \text{ and } L \leq \sigma^2\big) \leq \exp\left(-\frac{\alpha^2}{2\sigma^2 + 2M\alpha}\right)$$

3 Load-balancing using experts

In this section we formally define the generalized load-balancing problem and present our online algorithm for it. In the next section, we will use this algorithm to solve packing/covering LPs given an estimate of the optimal value.

3.1 Definitions: offline and online instances

An instance of the *offline* version of the *generalized load-balancing problem* is a set of matrices $\{A^t\}_t$, each in $\mathbb{R}^{m \times k}$. The goal is to find vectors $p^1, p^2, \ldots, p^n \in \Delta^k$ to minimize $\|\sum_{t=1}^n A^t p^t\|_{\max}$, the load of the most-loaded machine. We use λ^* throughout to denote the offline optimum value. Notice that, by adding a column of zeroes, we can get an instance where one option is to do nothing, i.e., we can change the problem to minimizing $\|\sum_{t=1}^n A^t p^t\|_{\max}$ over $p^t \in \blacktriangle^k$.

In the *online* version of this problem, we consider the random permutation model. Now the number of time steps n is the only information known upfront. Let $\mathbf{A}^1, \mathbf{A}^2, \ldots, \mathbf{A}^n$ be matrices sampled from the set $\{A^1, \ldots, A^n\}$ uniformly *without replacement*. At time step t, the algorithm has seen matrices $\mathbf{A}^1, \ldots, \mathbf{A}^t$ and must decide on a (random) vector $\mathbf{p}^t \in \Delta$; the randomness is both due to the random matrices seen so far and the internal coin flips of the algorithm (if any). The goal is to minimize $\|\sum_{t=1}^n \mathbf{A}^t \mathbf{p}^t\|_{\max}$. The vectors $\{\mathbf{p}^t\}_t$ output by the algorithm are called the *online solution* for the online instance $\{\mathbf{A}^t\}_t$ corresponding to the offline instance $\{A^t\}_t$.

Ideally we want guarantees that for any given $\varepsilon, \delta > 0$, we get $\|\sum_{t=1}^n \mathbf{A}^t \mathbf{p}^t\|_{\max} \leq (1+\varepsilon)\lambda^*$ with probability $(1-\delta)$. Due to the online nature of the problem, such a guarantee is only conceivable if λ^* is "large" compared to the maximum size of the loads. In the rest of this section we show how to get the best possible such guarantee.

3.2 Well-bounded instances

While instances arising from scheduling-type applications usually consist solely of positive loads, our reduction used for packing/covering LPs generates instances where some entries A^t_{ij} are also negative.

In order to get good guarantees, we need to control how negative these entries can be. Loosely speaking, while the entries of the load matrices A^t can be in the symmetric interval $[-M, M]$, for each machine we require that either the loads are mostly positive on all time steps (with any negative loads being tiny), or they are mostly negative on all time steps.

Definition 3.1 (Well-bounded instance)**.** *For $M, \gamma \geq 0$, an instance $A^1, \ldots, A^n$ of the load-balancing problem is* (M, γ)-well-bounded *if $A^t_{ij} \in [-M, M]$ for all i, j, t, and moreover for each $i \in [m]$ we have: either $A^t_{ij} \geq -\frac{\gamma\lambda^*}{n}$ for all j and $t \in [n]$, or $A^t_{ij} \leq \frac{\gamma\lambda^*}{n}$ for all j and $t \in [n]$.*

In particular, this is satisfied with $\gamma = 0$ if the A^t's are non-negative. The reader can think of γ as a small constant, say one. The main motivation behind this definition is that it allows us to control the variation of random processes defined over $\{A^t\}_t$ in a way similar to how we obtained Corollary 2.4.

Lemma 3.2. *Suppose $\{A^t\}_{t=1}^n$ is an (M, γ)-well-bounded instance for some $M, \gamma \geq 0$, and consider $\bar{p}^1, \ldots, \bar{p}^n \in \Delta^k$. Let the sequence $\mathbf{A}^1\bar{\mathbf{p}}^1, \ldots \mathbf{A}^n\bar{\mathbf{p}}^n$ be sampled without replacement from the set $\{A^t\bar{p}^t\}_t$. Then for every event $\mathcal{E}$ and for every $i \in [m]$ and $t \in [n]$,*

$$\mathbb{E}\left[\left|\mathbf{A}^t_i\bar{\mathbf{p}}^t - \mathbb{E}[\mathbf{A}^t_i\bar{\mathbf{p}}^t \,|\, \mathcal{E}]\right| \;\middle|\; \mathcal{E}\right] \leq \frac{2\gamma\lambda^*}{n} + 2\left|\mathbb{E}\left[\mathbf{A}^t_i\bar{\mathbf{p}}^t \,|\, \mathcal{E}\right]\right|.$$

Proof. Fix $i \in [m]$. Using the well-boundedness condition, define $\bar{\mathbf{v}}^t_i$ as follows: if $A^t_i\bar{p}^t \geq -\frac{\gamma\lambda^*}{n}$ for all t, set $\bar{\mathbf{v}}^t_i = \mathbf{A}^t_i\bar{\mathbf{p}}^t + \frac{\gamma\lambda^*}{n}$ for all t; else set $\bar{\mathbf{v}}^t_i = \frac{\gamma\lambda^*}{n} - \mathbf{A}^t_i\bar{\mathbf{p}}^t$. By definition for all t we have $\bar{\mathbf{v}}^t_i \geq 0$, and since we did a uniform shift

$$\mathbb{E}\left[\left|\mathbf{A}^t_i\bar{\mathbf{p}}^t - \mathbb{E}[\mathbf{A}^t_i\bar{\mathbf{p}}^t \,|\, \mathcal{E}]\right| \;\middle|\; \mathcal{E}\right] = \mathbb{E}\left[\left|\bar{\mathbf{v}}^t_i - \mathbb{E}[\bar{\mathbf{v}}^t_i \,|\, \mathcal{E}]\right| \;\middle|\; \mathcal{E}\right] \leq \mathbb{E}\left[|\bar{\mathbf{v}}^t_i| \;\middle|\; \mathcal{E}\right] + \left|\mathbb{E}\left[\bar{\mathbf{v}}^t_i \,|\, \mathcal{E}\right]\right| = 2\mathbb{E}\left[\bar{\mathbf{v}}^t_i \,|\, \mathcal{E}\right].$$

But it is easy to see that $\mathbb{E}[\bar{\mathbf{v}}^t_i \,|\, \mathcal{E}] \leq \frac{\gamma\lambda^*}{n} + \left|\mathbb{E}[\mathbf{A}^t_i\bar{\mathbf{p}}^t \,|\, \mathcal{E}]\right|$, which concludes the proof. □

3.3 The expertLB algorithm and its guarantee

Given an online instance $\{\mathbf{A}^t\}_t$ to the generalized load-balancing problem and values n, M, and ε, the following algorithm **expertLB** (for "expert load-balancing") runs a primal greedy strategy and for the dual **any** prediction-with-experts algorithm, restarting at timestep $n/2$.

Algorithm 1 expertLB

for each time t **do**
  let $\mathbf{A}^t$ be the arriving random item
  primal step: choose $\mathbf{p}^t \in \Delta^k$ to minimize $\left\langle \mathbf{w}^t\mathbf{A}^t, \mathbf{p}^t \right\rangle$ ▷ *play best-response*
  dual step: feed $\mathbf{A}^t\mathbf{p}^t$ to a prediction-with-experts algorithm (whose current state is $\mathbf{w}^t$)
    let $\mathbf{w}^{t+1}$ be the resulting state of the algorithm.
  if $t = n/2$ **then**
    restart the state of the prediction-with-experts algorithm ▷ *restart algorithm at half-time*
  end if
end for

Note the simplicity of the algorithm: it is perhaps the "natural" algorithm, once we decide to reduce load-balancing to a prediction-with-experts algorithm. The only non-intuitive step is perhaps that we restart the process at time $n/2$. This ensures that at each step, the current state depends on at most $n/2 - 1$ of the random choices from the $A^1, \ldots, A^n$, and hence the next random item still has "enough

randomness" for our analysis to go through[4]. Also, another interpretation of the algorithm is that the dual vectors $\mathbf{w}^t \in \Delta^m$ are providing (adaptive) aggregations of the m machines into a single one, in which case the best action is just to minimize the load on this machine (which is exactly what our best-response primal step is doing).

The following theorem shows if the offline optimal load λ^* is not too small, then **expertLB** obtains value almost equal to λ^* up to the regret of the prediction-with-experts algorithm used.

Theorem 3.3 (Load-balancing guarantee)**.** *Suppose $\{A^t\}_t$ is (M,γ)-well-bounded load-balancing instance for $\gamma \geq 1$. Let λ^* be its optimal load and suppose $\varepsilon \leq \varepsilon_1 := \frac{1}{84}$ and $\delta \in (0,\varepsilon]$ are such that $\lambda^* \geq \frac{3M\log(m/\delta)}{\varepsilon^2}$. If the algorithm **expertLB** is run with **any** prediction-with-experts algorithm with (α, R)-regret, then it returns a solution $\{\mathbf{p}^t\}_t$ such that with probability at least $1-\delta$*

$$\left\| \sum_t \mathbf{A}^t\mathbf{p}^t - \alpha \sum_t |\mathbf{A}^t\mathbf{p}^t| \right\|_{\max} \leq \lambda^*(1 + c_1\gamma\varepsilon) + 2R,$$

*for a universal constant c_1. In particular, if the **MWU** algorithm from Theorem 2.2 is used with learning parameter ε, the guarantee becomes (for a slightly larger constant c_1)*

$$\left\| \sum_t \mathbf{A}^t\mathbf{p}^t - \varepsilon \sum_t |\mathbf{A}^t\mathbf{p}^t| \right\|_{\max} \leq \lambda^*(1 + c_1\gamma\varepsilon).$$

The case when the load matrices are non-negative is of independent interest (generalizing scheduling on unrelated machines); in this case well-boundedness means $A^t_{ij} \in [0, M]$ and we can set $\gamma = 1$ to get:

Corollary 3.4 (Positive loads guarantee)**.** *Suppose $\{A^t\}_t$ is a load-balancing instance with $A^t_{ij} \in [0, M]$, and optimal load λ^*. Suppose $\varepsilon \leq \varepsilon_1 := \frac{1}{84}$ and $\delta \in (0,\varepsilon]$ satisfy $\lambda^* \geq \frac{3M\log(m/\delta)}{\varepsilon^2}$. Given n, M and ε, the algorithm **expertLB** (with the **MWU** algorithm) finds an online solution $\{\mathbf{p}^t\}_t$ such that $\|\sum_t \mathbf{A}^t\mathbf{p}^t\|_\infty \leq (1 + O(\varepsilon))\lambda^*$ with probability at least $1-\delta$.*

In the rest of §3 we prove Theorem 3.3. We first show in §3.4 that the maximum load is close in expectation to λ^*. We then extend these ideas to prove the high-probability guarantee in §3.5.

3.4 The guarantee in expectation

Let $p^{*1}, \ldots, p^{*n}$ be the optimal solution for the offline instance. We see this as a mapping from matrix A^t to solution p^{*t} (say $\phi(A^t) = p^{*t}$); this way, define $\mathbf{p}^{*t}$ as the random solution with respect to the random matrix $\mathbf{A}^t$ (i.e. $\mathbf{p}^{*t} = \phi(\mathbf{A}^t)$).

To simplify the notation, let $\mathbf{o}^t := \mathbf{A}^t\mathbf{p}^t$ denote the load vector incurred at step t by our algorithm. Also, for an integer ℓ, we use $\mathbf{A}^{\leq \ell}$ to denote the sequence $\mathbf{A}^1, \ldots, \mathbf{A}^\ell$, and similarly for other objects.

Theorem 3.3 seeks to bound $\left\|\sum_t \mathbf{o}^t - \alpha\sum_t |\mathbf{o}^t|\right\|_{\max}$. By the exchangeability of sampling without replacement, $(\mathbf{A}^1, \ldots, \mathbf{A}^{n/2})$ has the same distribution as $(\mathbf{A}^{n/2+1}, \ldots, \mathbf{A}^n)$. This implies that our algorithm, with its half-time reset, behaves in the same way in the two halves of the process, namely $(\mathbf{A}^{\leq n/2}, \mathbf{p}^{\leq n/2})$ has the same distribution as $(\mathbf{A}^{>n/2}, \mathbf{p}^{>n/2})$; thus, it suffices to analyze the first half.

Fact 3.5 (Suffices to analyze first half)**.** *The random variables $\left\| \sum_{t\leq n/2} \mathbf{o}^t - \alpha \sum_{t\leq n/2} |\mathbf{o}^t| \right\|_{\max}$ and $\left\| \sum_{t>n/2} \mathbf{o}^t - \alpha \sum_{t>n/2} |\mathbf{o}^t| \right\|_{\max}$ have the same distribution.*

[4]The formal reason why we need this restart is that Lemma 2.5 (because of Theorem A.1) can only handle up to half of the random sequence.

Now we need that the computed dual solutions $\mathbf{w}^t$ capture our (non-linear) total load, i.e., that $\sum_t \langle \mathbf{w}^t, \mathbf{o}^t \rangle \approx \| \sum_t \mathbf{o}^t \|_{\max}$; but this follows directly from the fact we used an (α, R)-regret prediction-with-experts algorithm.

Fact 3.6 (Dual captures load)**.** *For every scenario we have*

$$\sum_{t=1}^{n/2} \langle \mathbf{w}^t, \mathbf{o}^t \rangle \;\; \geq \;\; \Big\| \sum_{t \leq n/2} \mathbf{o}^t - \alpha \sum_{t \leq n/2} |\mathbf{o}^t| \Big\|_{\max} - R.$$

Furthermore, let $\mathbf{o}^{*t} := \mathbf{A}^t \mathbf{p}^{*t}$ denote the load incurred by the optimal solution $\mathbf{p}^*$ at step t. By our primal greedy choice of the $\mathbf{p}^t$'s, we directly have the following.

Fact 3.7 (Optimality of algorithm wrt duals)**.**

$$\sum_{t=1}^{n/2} \langle \mathbf{w}^t, \mathbf{o}^t \rangle = \sum_{t=1}^{n/2} \mathbf{w}^t \mathbf{A}^t \mathbf{p}^t \leq \sum_{t=1}^{n/2} \mathbf{w}^t \mathbf{A}^t \mathbf{p}^{*t} = \sum_{t=1}^{n/2} \langle \mathbf{w}^t, \mathbf{o}^{*t} \rangle.$$

Together these give $\| \sum_{t \leq n/2} \mathbf{o}^t - \alpha \sum_{t \leq n/2} |\mathbf{o}^t| \, \|_{\max} \leq R + \sum_{t=1}^{n/2} \langle \mathbf{w}^t, \mathbf{o}^{*t} \rangle$ and so in order to upper bound our load in expectation it suffices to show that

$$\mathbb{E}\Big[\sum_{t \leq n/2} \langle \mathbf{w}^t, \mathbf{o}^{*t} \rangle \Big] \lesssim \sum_{t \leq n/2} \langle \mathbb{E}[\mathbf{w}^t], \mathbb{E}[\mathbf{o}^{*t}] \rangle \leq \Big\| \mathbb{E}\Big[\sum_{t \leq n/2} \mathbf{o}^{*t} \Big] \Big\|_{\max} = \frac{\lambda^*}{2}. \tag{3.1}$$

Notice that these are easy implications if we were sampling *with replacement*, since in that case $\mathbf{w}^t$ and $\mathbf{o}^{*t}$ are independent. For the random permutation model we need to work harder because they are dependent: $\mathbf{w}^t$ is determined by the history $\mathbf{A}^{<t}$ and the optimal solution's load $\mathbf{o}^{*t}$ at time t depends on the history $\mathbf{A}^{<t}$ (since items that have occurred in the history cannot occur at time t).

However, this dependence of $\mathbf{o}^{*t}$ on the history $\mathbf{A}^{<t}$ should be quite small. In fact this is the main tool for our analysis: we show that $\mathbb{E}[\mathbf{o}_i^{*t} \,|\, \mathbf{A}^{<t}] \approx \mathbb{E}\, \mathbf{o}_i^{*t}$ happens for most histories, for all t, i. More precisely, we show that for each i, with high probability $\mathbb{E}[\mathbf{o}_i^{*t} \,|\, \mathbf{A}^{<t}] = \mathbb{E}\, \mathbf{o}_i^{*t} \pm O(\varepsilon) \cdot (|\mathbb{E}\, \mathbf{o}_i^{*t}| + \frac{\gamma \lambda^*}{n})$ *for all* $t \leq n/2$ *simultaneously*. Note that applying Bernstein's inequality to each $\mathbf{o}_i^{*t}$ and taking a union bound over the t's would only give $\mathbb{E}[\mathbf{o}_i^{*t} \,|\, \mathbf{A}^{<t}] = \mathbb{E}\, \mathbf{o}_i^{*t} \pm O(\varepsilon \log n) \cdot (|\mathbb{E}\, \mathbf{o}_i^{*t}| + \frac{\gamma \lambda^*}{n})$, with an extra $\log n$ factor. To avoid this we employ a *maximal* version of Bernstein's inequality. To simplify the notation, we use the shorthand $\mathbb{E}_t \mathbf{X} := \mathbb{E}[\mathbf{X} \,|\, \mathbf{A}^{<t}]$ for expectations conditioned on the history before time t.

Lemma 3.8. *Consider $i \in [m]$. Then with probability at least $1 - \frac{\delta}{m^2}$ we have $\mathbb{E}_t[\mathbf{o}_i^{*t}] \in \mathbb{E}\, \mathbf{o}_i^{*t} \pm 80\varepsilon(|\mathbb{E}\, \mathbf{o}_i^{*t}| + \frac{\gamma \lambda^*}{n})$ for all $t \leq \frac{n}{2}$ simultaneously.*

Proof. Let $\mu = \frac{1}{n} \sum_{t \leq n} A_i^t p^{*t} \leq \frac{\lambda^*}{n}$ be the expected value of $\mathbf{o}_i^{*t}$, which is independent of t. Notice that in every scenario $\sum_{j=1}^{n} \mathbf{o}_i^{*j} - \sum_{j<t} \mathbf{o}_i^{*j}$ gives the total i^{th} load from items that have not shown up at times $j < t$. Therefore, the conditional expected i^{th} load at time t is the average over these remaining items:

$$\mathbb{E}_t[\mathbf{o}_i^{*t}] = \frac{\sum_{j=1}^{n} \mathbf{o}_i^{*j} - \sum_{j<t} \mathbf{o}_i^{*j}}{n - (t-1)} = \frac{n\mu - \sum_{j<t} \mathbf{o}_i^{*j}}{n - (t-1)}. \tag{3.2}$$

Thus, to control $\mathbb{E}_t[\mathbf{o}_i^{*t}]$ it suffices to control the sum $\sum_{j<t} \mathbf{o}_i^{*j}$, which should be highly concentrated around the mean $(t-1)\mu$.

For that, let $\alpha = n|\mu| + \gamma\lambda^*$; we want to use the maximal Bernstein's inequality for sampling without replacement (Lemma 2.5) to bound the maximum deviation $\max_{t \leq n/2} |\sum_{j<t} \mathbf{o}_i^{*j} - t\mu| \leq 40\varepsilon\alpha$ with

probability at least $1-\frac{\delta}{m^2}$. For that, we need to bound the variance $V := \frac{1}{n}\sum_{t\le n}(A_i^t p^{*t}-\mu)^2 = \mathbb{E}(\mathbf{o}_i^{*t}-\mathbb{E}\,\mathbf{o}_i^{*t})^2$. Since $\mathbf{o}_i^{*t}\in[-M,M]$, we have that $V \le M\cdot\mathbb{E}\big[\big|\mathbf{o}_i^{*t}-\mathbb{E}\,\mathbf{o}_i^{*t}\big|\big]$. Then employing Lemma 3.2 to bound this absolute value we get $V \le M(2|\mu|+\frac{2\gamma\lambda^*}{n}) = 2M\alpha/n$. Thus, applying the maximal Bernstein's inequality Lemma 2.5 we obtain

$$\Pr\left(\max_{t\le n/2}|\textstyle\sum_{j\le t}\mathbf{o}_i^{*j}-t\mu|\le 40\varepsilon\alpha\right) \le 30\exp\left(-\left(\frac{40\varepsilon}{24}\right)^2\frac{\alpha^2}{nV+40\varepsilon\alpha(M/24)}\right)$$
$$\le 30\exp\left(-\left(\frac{40\varepsilon}{24}\right)^2\frac{\alpha}{(2+40\varepsilon/24)M}\right) \le 30\exp\left(-\frac{\varepsilon^2\gamma\lambda^*}{M}\right) \le \frac{\delta}{m^2},$$

where we used $\delta\le\varepsilon\le\frac{1}{80}$, $\gamma\ge 1$, and $\lambda^*\ge\frac{3M\log(m/\delta)}{\varepsilon^2}$.

Whenever this event holds equation (3.2) gives that for any $t\le n/2$

$$\mathbb{E}_t[\mathbf{o}_i^{*t}] \in \frac{(n-(t-1))\,\mu\pm 40\varepsilon\alpha}{n-(t-1)} = \mu\pm\frac{40\varepsilon\alpha}{n-(t-1)} \in \mu\pm\frac{80\varepsilon\alpha}{n},$$

where the last inequality uses $t\le n/2$. The lemma then follows from the definition of α. □

Using this lemma it is not difficult to show that inequality (3.1) approximately holds. We sketch how this can be done in order to illustrate that little of the structure of the problem is actually used (and thus might carry over to other settings). This inequality plus Facts 3.5–3.7 imply that the bound in Theorem 3.3 holds in expectation. (We remark that this guarantee in expectation does not require Lemma 3.8 to hold simultaneously for all $t\le n/2$; this will only be required below for our guarantee with high probability.)

Lemma 3.9. *We have* $\mathbb{E}\Big[\sum_{t\le n/2}\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\Big]\le\frac{\lambda^*}{2}+O(\varepsilon\gamma\lambda^*)$.

Proof sketch. Let $\mathcal{E}$ denote the event that the approximation from Lemma 3.8 holds for every $i\in[m]$, which from a union bound then holds with probability at least $1-\frac{\delta}{m}$. Let us split the expectation in the left-hand side of the lemma based on $\mathcal{E}$:

$$\mathbb{E}\Big[\sum_{t\le n/2}\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\Big] = \mathbb{E}\Big[\sum_{t\le n/2}\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\;\Big|\;\mathcal{E}\Big]\Pr(\mathcal{E}) + \mathbb{E}\Big[\sum_{t\le n/2}\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\;\Big|\;\overline{\mathcal{E}}\Big]\Pr(\overline{\mathcal{E}}). \tag{3.3}$$

To upper bound the first term we further split the expectation by conditioning on the history up to time $t-1$ and recalling that $\mathbf{w}^t$ is defined by it:

$$\mathbb{E}\Big[\sum_{t\le n/2}\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\;\Big|\;\mathcal{E}\Big] = \sum_{t\le n/2}\mathbb{E}\Big[\mathbb{E}\Big[\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\,|\,\mathbf{A}^{<t},\mathcal{E}\Big]\;\Big|\;\mathcal{E}\Big] = \sum_{t\le n/2}\mathbb{E}\Big[\big\langle\mathbf{w}^t,\mathbb{E}[\mathbf{o}^{*t}\,|\,\mathbf{A}^{<t},\mathcal{E}]\big\rangle\;\Big|\;\mathcal{E}\Big]. \tag{3.4}$$

Recalling the guarantee of Lemma 3.8 for $\mathcal{E}$, we can bounds the innermost expectation by

$$\mathbb{E}[\mathbf{o}^{*t}\,|\,\mathbf{A}^{<t},\mathcal{E}] \le \mathbb{E}\,\mathbf{o}^{*t}+80\varepsilon|\mathbb{E}\,\mathbf{o}^{*t}|+O(\varepsilon\gamma)\frac{\lambda^*}{n}\mathbb{1} \le \max\{(1+80\varepsilon)\mathbb{E}\,\mathbf{o}^{*t},0\}+O(\varepsilon\gamma)\frac{\lambda^*}{n}\mathbb{1} \le (1+O(\varepsilon\gamma))\frac{\lambda^*}{n}\mathbb{1},$$

where the second inequality uses that $\varepsilon\le\frac{1}{80}$ and the last one uses the optimality $\mathbf{o}^{*t}\le(\lambda^*/n)\mathbb{1}$. Plugging this bound onto inequality (3.4) gives that we can upper bound the first term in equation (3.3) by $(1+O(\varepsilon\gamma))\frac{\lambda^*}{2}$.

To upper bound the second term in equation (3.3), using the well-boundedness of the instance it is not hard to show that in every scenario the sum of optimal loads satisfies $\sum_{t\le n/2}\mathbf{o}_i^{*t}\le O(\gamma)\lambda^*$, and thus $\sum_{t\le n/2}\big\langle\mathbf{w}^t,\mathbf{o}^{*t}\big\rangle\le O(\gamma)\,m\lambda^*$. Since $\bar{\mathcal{E}}$ only holds with probability $\frac{\delta}{m}\le\frac{\varepsilon}{m}$, we can upper bound the second term in equation (3.3) by $O(\varepsilon\gamma)\,\lambda^*$.

Putting these bounds together concludes the proof of the lemma. □

3.5 The guarantee with high probability

We now show that Lemma 3.9 actually holds with high probability.

Lemma 3.10. *With probability at least $1-3\delta$ we have*

$$\textstyle\sum_{t\le n/2}\langle \mathbf{w}^t, \mathbf{o}^{*t}\rangle \le (1+168\varepsilon\gamma)\frac{\lambda^*}{2}.$$

Again, if we were sampling *with replacement*, then it would not be difficult to show that $\big(\langle \mathbf{w}^t, \mathbf{o}^{*t}\rangle - \langle \mathbf{w}^t, \mathbb{E}\,\mathbf{o}^{*t}\rangle\big)_t$ is a martingale difference sequence, and consequently martingale concentration inequalities would imply $\sum_{t\le n/2}\langle \mathbf{w}^t, \mathbf{o}^{*t}\rangle \lesssim \sum_{t\le n/2}\langle \mathbf{w}^t, \mathbb{E}\,\mathbf{o}^{*t}\rangle \le \sum_{t\le n/2}\|\mathbb{E}\,\mathbf{o}^{*t}\|_{\max} \lesssim \frac{\lambda^*}{2}$ with good probability. For our case of sampling without replacement, we proceed similarly, but using the martingale difference sequence $\big(\langle \mathbf{w}^t, \mathbf{o}^{*t}\rangle - \langle \mathbf{w}^t, \mathbb{E}_t[\mathbf{o}^{*t}]\rangle\big)_t$ (notice the conditioning on the history), and then using Lemma 3.8 to get that with good probability for all t we have $\mathbb{E}_t[\mathbf{o}^{*t}] \approx \mathbb{E}[\mathbf{o}^{*t}]$.

So define the random variable $\mathbf{Y}_t := \langle \mathbf{w}^t, \mathbf{o}^{*t} - \mathbb{E}_t[\mathbf{o}^{*t}]\rangle$. Since $\mathbf{w}^t$ is determined by $\mathbf{A}^{<t}$, we have that $\mathbb{E}_t[\mathbf{Y}_t]=0$ and hence $(\mathbf{Y}_t)_t$ is indeed a martingale difference sequence (adapted to $(\mathbf{A}^t)_t$). In order to obtain concentration for this martingale, we first need to control its conditional variance, which is accomplished in the next lemma.

Lemma 3.11. *With probability at least $1-\frac{\delta}{m}$ we have*

$$\sum_{t\le n/2}\mathrm{Var}(\mathbf{Y}_t \mid \mathbf{A}^{<t}) \le 4M\cdot\sum_{t\le n/2}\big\langle \mathbf{w}^t, \big|\mathbb{E}\,\mathbf{o}^{*t}\big|\big\rangle + 4M\gamma\lambda^*.$$

Proof. Since $\mathbf{o}_i^{*t}\in[-M,M]$, we have $\mathbf{Y}_t\in[-2M,2M]$ and hence $\mathbf{Y}_t^2 \le 2M|\mathbf{Y}_t|$. Since $\mathbb{E}_t[\mathbf{Y}_t]=0$, for all the histories we have

$$\begin{aligned}\mathrm{Var}(\mathbf{Y}_t \mid \mathbf{A}^{<t}) = \mathbb{E}_t[\mathbf{Y}_t^2] &\le 2M\cdot\mathbb{E}_t|\mathbf{Y}_t| = 2M\cdot\mathbb{E}_t\Big[\Big|\sum_i \mathbf{w}_i^t\big(\mathbf{o}_i^{*t}-\mathbb{E}_t[\mathbf{o}_i^{*t}]\big)\Big|\Big]\\ &\le 2M\sum_i \mathbf{w}_i^t\,\mathbb{E}_t\big[\big|\mathbf{o}_i^{*t}-\mathbb{E}_t[\mathbf{o}_i^{*t}]\big|\big] \overset{Lemma\ 3.2}{\le} 4M\,\big\langle \mathbf{w}^t, \big|\mathbb{E}_t[\mathbf{o}^{*t}]\big|\big\rangle + \frac{4M\gamma\lambda^*}{n},\end{aligned}$$

where the last inequality uses $\mathbf{w}^t\in\Delta^m$.

To conclude the proof, we need the approximation $\mathbb{E}[\mathbf{o}^{*t}]\approx\mathbb{E}_t[\mathbf{o}^{*t}]$. For that, let $\mathcal{E}$ be the event that the guarantees from Lemma 3.8 hold for all $i\in[m]$; by a union bound, we have that $\mathcal{E}$ happens with probability at least $1-\frac{\delta}{m}$. Under the event $\mathcal{E}$, we have that for all $t\le n/2$

$$|\mathbb{E}_t\,\mathbf{o}_i^{*t}| \le |\mathbb{E}\,\mathbf{o}_i^{*t}| + 80\varepsilon|\mathbb{E}\,\mathbf{o}_i^{*t}| + \frac{80\varepsilon\gamma\lambda^*}{n} \le 2|\mathbb{E}\,\mathbf{o}_i^{*t}| + \frac{\gamma\lambda^*}{n},$$

where the last inequality uses $\varepsilon\le\frac{1}{80}$. Thus, under the event $\mathcal{E}$ we have $\mathrm{Var}(\mathbf{Y}_t\mid\mathbf{A}^{<t}) \le 8M\big\langle \mathbf{w}^t, |\mathbb{E}\,\mathbf{o}^{*t}|\big\rangle + \frac{8M\gamma\lambda^*}{n}$. Adding over all $t\le n/2$ then gives the lemma. ☐

Now we can use our martingale to prove Lemma 3.10.

Proof of Lemma 3.10. We first employ Freedman's inequality to the martingale $(\mathbf{Y}_t)_t$. For that, let $\alpha := 4\sum_{t\le n/2}\big\langle \mathbf{w}^t, |\mathbb{E}\,\mathbf{o}^{*t}|\big\rangle + 4\gamma\lambda^*$, so that $M\alpha$ upper bounds the conditional variance $\sigma^2 := \sum_{t\le n/2}\mathrm{Var}(\mathbf{Y}_t\mid\mathbf{A}^{<t})$ with probability $1-\frac{\delta}{m}$. Then from Freedman's inequality (Theorem 2.6) we get that

$$\Pr\Big(\sum_{t\le n/2}\mathbf{Y}_t > \varepsilon\alpha \text{ and } \sigma^2\le M\alpha\Big) \le \exp\Big(-\frac{\varepsilon^2\alpha^2}{4M\alpha}\Big) \le \exp\Big(-\frac{\varepsilon^2\gamma\lambda^*}{M}\Big) \le \frac{\delta}{m},$$

where the last inequality uses $\gamma \geq 1$ and $\lambda^* \geq \frac{3M\log(m/\delta)}{\varepsilon^2}$. Since $\Pr(\sigma^2 \leq M\alpha) \geq 1 - \frac{\delta}{m}$, a union bound gives that with probability at least $1 - \frac{2\delta}{m}$ we have $\sum_{t\leq n/2} \mathbf{Y}_t \leq \varepsilon\alpha$. Expanding the definition of $\mathbf{Y}$ and α, this is equivalent to

$$\sum_{t\leq n/2} \left\langle \mathbf{w}^t, \mathbf{o}^{*t} \right\rangle \leq \sum_{t\leq n/2} \left\langle \mathbf{w}^t, \mathbb{E}_t\, \mathbf{o}^{*t} + 4\varepsilon|\mathbb{E}\, \mathbf{o}^{*t}| \right\rangle + 4\varepsilon\gamma\lambda^*. \tag{3.5}$$

To upper bound the first term in the right-hand side, employing Lemma 3.8 to each $i \in [m]$ we have that with probability at least $1 - \frac{\delta}{m}$ the following holds for all $t \leq n/2$:

$$\mathbb{E}_t\, \mathbf{o}^{*t} + 4\varepsilon|\mathbb{E}\, \mathbf{o}^{*t}| \leq \mathbb{E}\, \mathbf{o}^{*t} + 84\varepsilon|\mathbb{E}\, \mathbf{o}^{*t}| + \frac{80\varepsilon\gamma\lambda^*}{n}\mathbb{1} \leq \max\{(1+84\varepsilon)\mathbb{E}\, \mathbf{o}^{*t}, 0\} + \frac{80\varepsilon\gamma\lambda^*}{n}\mathbb{1} \leq (1+164\varepsilon\gamma)\frac{\lambda^*}{n},$$

where the second inequality follows from the fact $\varepsilon \leq \frac{1}{84}$ and $\mathbb{E}\, \mathbf{o}^{*t} \leq (\lambda^*/n)\mathbb{1}$. Plugging this bound on inequality (3.5) and taking a union bound, we have that with probability at least $1 - \frac{3\delta}{m}$

$$\sum_{t\leq n/2} \left\langle \mathbf{w}^t, \mathbf{o}^{*t} \right\rangle \leq (1+168\varepsilon\gamma)\frac{\lambda^*}{2}.$$

This concludes the proof. □

Proof of Theorem 3.3. From triangle inequality we have

$$\left\| \sum_t \mathbf{o}^t - \alpha \sum_t |\mathbf{o}^t| \right\|_{\max} \leq \left\| \sum_{t\leq n/2} \mathbf{o}^t - \alpha \sum_{t\leq n/2} |\mathbf{o}^t| \right\|_{\max} + \left\| \sum_{t>n/2} \mathbf{o}^t - \alpha \sum_{t>n/2} |\mathbf{o}^t| \right\|_{\max}.$$

Using Facts 3.6 and 3.7 and Lemma 3.10, with probability at least $1 - 3\delta$ the first term in the right-hand side is at most $(1+168\varepsilon\gamma)\frac{\lambda^*}{2} + R$. Since from Fact 3.5 the same bound holds for the second term of the right-hand side, a union bound concludes the proof of the theorem with constant $c_1 = 168$. □

4 Solving packing-covering multi-choice LPs given estimate of OPT

This section shows how to solve packing/covering multiple-choice (PCMC) LPs via a simple reduction to the generalized load-balancing problem from the previous section, provided an estimate of the optimal value is available. In the following sections we show how to get such an estimate.

From §1.1, recall that a PCMC LP has the form:

$$\begin{aligned} \max \ & \textstyle\sum_{t=1}^n \pi^t x^t && \text{(PCMC)} \\ st \ & \textstyle\sum_{t=1}^n A^t x^t \leq b && (4.6) \\ & \textstyle\sum_{t=1}^n C^t x^t \geq d && (4.7) \\ & x^t \in \blacktriangle^k \quad \forall t \in [n], && (4.8) \end{aligned}$$

where all the data π^t, A^t, C^t, b, d is **non-negative** and $\blacktriangle^k$ denotes the "full simplex" $\{x \in [0,1]^k : \sum_j x_j \leq 1\}$. We use m_p to denote the number packing constraints (4.6) and m_c to denote the number of covering constraints (4.7), so our matrices have dimensions $A^t \in \mathbb{R}_+^{m_p \times k}$, and $C^t \in \mathbb{R}_+^{m_c \times k}$. Given a PCMC LP $\mathcal{L}$, we use $\mathsf{opt}(\mathcal{L})$ to denote its optimal value; we use simply opt when the LP is clear from context.

In the *online* version of the problem there is a fixed PCMC LP $\mathcal{L}$ whose blocks $\{(\pi^t, A^t, C^t)\}_t$ are presented one-by-one in random order. More precisely, the number of time steps n and the right-hand sides b, d are known upfront, and the triples $(\boldsymbol{\pi}^1, \mathbf{A}^1, \mathbf{C}^1), \ldots, (\boldsymbol{\pi}^n, \mathbf{A}^n, \mathbf{C}^n)$ are sampled from

$\{(\pi^t, A^t, C^t)\}_t$ uniformly *without replacement*. Then define the randomly permuted LP $\boldsymbol{\mathcal{L}}$

$$\boldsymbol{\mathcal{L}} = \max\left\{ \sum_t \boldsymbol{\pi}^t x^t : \sum_t \mathbf{A}^t x^t \le b,\ \sum_t \mathbf{C}^t x^t \ge d,\ x^t \in \blacktriangle^k \ \ \forall t \right\} \tag{4.9}$$

At time step t, the algorithm computes a (random) vector $\mathbf{x}^t \in \blacktriangle^k$ based on the information seen up to time t, i.e., $(\boldsymbol{\pi}^1, \mathbf{A}^1, \mathbf{C}^1), \ldots, (\boldsymbol{\pi}^t, \mathbf{A}^t, \mathbf{C}^t)$, plus n and b, d. We call the sequence $\{\mathbf{x}^t\}_t$ an *online solution*.

We say that an online solution $\{\mathbf{x}^t\}_t$ is ε-*feasible* for $\boldsymbol{\mathcal{L}}$ if it satisfies the packing constraints (i.e., $\sum_{t\le n} \mathbf{A}^t \mathbf{x}^t \le b$) and *almost* satisfies the covering constraints (i.e., $\sum_{t\le n} \mathbf{C}^t \mathbf{x}^t \ge (1-\varepsilon)d$). The goal in the online PCMC LP problem is to obtain an online solution $\{\mathbf{x}^t\}_t$ which (with high probability) is ε-feasible, and obtains a reward $\sum_{t\le n} \boldsymbol{\pi}^t \mathbf{x}^t \ge (1 - O(\varepsilon))\mathsf{opt}(\boldsymbol{\mathcal{L}})$. Notice that $\mathsf{opt}(\boldsymbol{\mathcal{L}}) = \mathsf{opt}(\mathcal{L})$, so again we are comparing against the optimal offline solution.

4.1 The algorithm LPviaLB

We now give an algorithm **LPviaLB** that solves PCMC LPs via a natural reduction to the Load-balance problem. The idea is simple: The Load-balance problem can be thought as a congestion-minimization problem $\min\{\lambda : \sum_{t=1}^n A^t p^t \le \lambda \cdot \mathbb{1},\ p^t \in \Delta^k\ \forall t\}$, which is reminiscent of a *packing-only* LP where all right-hand sides are the same. To bring a PCMC LP into this form, we introduce the objective function as a covering constraint $\sum_{t=1}^n \pi^t x^t \ge \widehat{\mathsf{opt}}$ (thereby requiring an estimate of opt), multiply the covering constraints by -1 to obtain packing constraints, and shift/scale to ensure uniform right-hand sides. Observe that the reduction creates negative entries on the left-hand side, which is why we considered the Load-balance problem with possibly negative loads.

More precisely, given a PCMC LP $\mathcal{L}$ and an estimate $\widehat{\mathsf{opt}}$ of opt, define the matrices $H^1, \ldots, H^n$ with $k+1$ columns and $m_p + m_c + 1$ rows (indexed from 0 to $m := m_p + m_c$) as follows: the zeroth row of H^t equals the vector

$$H^t_{0,\star} := \left(\tfrac{2}{n} - \tfrac{\pi^t_1}{\widehat{\mathsf{opt}}}, \ldots, \tfrac{2}{n} - \tfrac{\pi^t_k}{\widehat{\mathsf{opt}}}, \tfrac{2}{n}\right);$$

for $i \in \{1, \ldots, m_p\}$, the i^{th} row of H^t is

$$H^t_{i,\star} := \left(\tfrac{a^t_{i1}}{b_i}, \ldots, \tfrac{a^t_{ik}}{b_i}, 0\right),$$

and for $i \in \{m_p + 1, \ldots, m_p + m_c\}$, the $(m_p + i)^{th}$ row of H^t is

$$H^t_{m_p+i,\star} := \left(\tfrac{2}{n} - \tfrac{c^t_{i1}}{d_i}, \ldots, \tfrac{2}{n} - \tfrac{c^t_{ik}}{d_i}, \tfrac{2}{n}\right).$$

The algorithm **LPviaLB** can be thought of as running in phases, where it constructs the instance $\{H^t\}_t$, feeds it to the algorithm **expertLB**, and outputs a scaling of the solution returned.

Algorithm 2 . LPviaLB

Input: Reals $\widehat{\mathsf{opt}}, \varepsilon, \delta > 0$, number n of items, and right-hand sides b and d.

Phase 1: Compute matrices $H^1, \ldots, H^n$ relative to the PCMC LP $\mathcal{L}$.

Phase 2: Run algorithm **expertLB** (with the **MWU** algorithm from Theorem 2.2) over the generalized load-balancing instance $\{H^t\}_t$ with parameters $M = \frac{2\varepsilon^2}{\log((m+1)/\delta)}$, $\varepsilon' = \sqrt{8}\varepsilon$ and δ (given), and let $\{\tilde{\mathbf{x}}^t\}_t$ be the returned solution.

Phase 3: Define $\widehat{\mathbf{x}}^t_j := \frac{(1-\varepsilon')}{(1+4c_1\varepsilon')}\tilde{\mathbf{x}}^t_j$ for $j = 1, \ldots, k$. Return $\{\widehat{\mathbf{x}}^t\}_t$ as the solution for the PCMC LP $\mathcal{L}$.

Some remarks are in order. Firstly, for simplicity we described the algorithm as running in phases, but it is easy to implement this in an online fashion. Indeed, when at time t the algorithm sees a random triple $(\boldsymbol{\pi}^t, \mathbf{A}^t, \mathbf{C}^t)$, it can construct the appropriate matrix $\mathbf{H}^t$, feed it to algorithm **expertLB** to get back the vector $\tilde{\mathbf{x}}^t$, and output the scaled version $\hat{\mathbf{x}}^t$. Secondly, we used the algorithm **expertLB** instantiated with the **MWU** algorithm, but this is for ease of exposition. We use the guarantees from Theorem 3.3 in a black-box fashion, and hence could use any algorithm for the prediction-with-experts problem (and obtaining different guarantees).

The guarantee of algorithm **LPviaLB** depends on the *generalized width* of the PCMC LP, defined as:

$$\min\left\{\min_{i,j,t}\frac{b_i}{a^t_{i,j}},\quad \min_{i,j,t}\frac{d_i}{c^t_{i,j}},\quad \min_{t,j}\frac{\mathsf{opt}}{\pi^t_j}\right\}. \tag{4.10}$$

The next theorem shows that algorithm **LPviaLB** with high probability returns an ε-feasible $(1-\varepsilon)$-optimal solution as long as the generalized width is $\Omega(\frac{\log(m/\delta)}{\varepsilon^2})$ and a $(1-\varepsilon)$-optimal under estimate $\widehat{\mathsf{opt}}$ is available.

.

Theorem 4.1. *Consider a feasible packing/covering multiple-choice LP $\mathcal{L}$. Consider $\varepsilon > 0$ at most a sufficiently small constant ε_0, and $\delta \in (0, \varepsilon]$, and assume that the generalized width of $\mathcal{L}$ is at least $\frac{\log((1+m_p+m_c)/\delta)}{\varepsilon^2}$. Then **LPviaLB** finds an online solution $\{\mathbf{x}^t\}$ that with probability at least $1-\delta$ is $(c_2\varepsilon)$-feasible for $\boldsymbol{\mathcal{L}}$ and has value $\sum_t \boldsymbol{\pi}^t\mathbf{x}^t \geq (1-c_3\varepsilon)\widehat{\mathsf{opt}}$, for constants $c_2, c_3 \geq 1$.*

4.2 Proof of Theorem 4.1

The idea is to first show that the reduction creates a load-balancing instance that is well-bounded, and whose optimal load is not very small. With these guarantees, we apply Theorem 3.3 from §3 to get a good load-balancing solution. Interpreting this as a solution to our PCMC LP will complete the proof.

Assume that all hypotheses in the statement of Theorem 4.1 hold. Let $\{H^t\}_t$ be the generalized load-balancing instance relative to the input PCMC LP $\mathcal{L}$, and let λ^* denote the optimal value of the instance $\{H^t\}_t$. We start by bounding λ^* by relating the solutions for the instance $\{H^t\}_t$ to the solutions for the LP $\mathcal{L}$.

Claim 4.2. $\frac{1}{2} \leq \lambda^* \leq 1$.

Proof. (Upper bound.) Let $\{x^t\}_t$ be an optimal solution for the PCMC LP $\mathcal{L}$, and for each t define $\tilde{x}^t = (x^t_1, \ldots, x^t_k, 0)$ by appending a zero at the end of x^t; it suffices to show that $\{\tilde{x}^t\}_t$ is a solution for the load-balancing instance $\{H^t\}_t$ with load at most 1.

The feasibility of $\{x^t\}_t$ guarantees that for $i \in [m_p]$ we have $\sum_t H^t_{i,\star}\tilde{x}^t = \frac{1}{b_i}\sum_t A^t_{i,\star}x^t \leq 1$, and for $i \in [m_c]$ we have $\sum_t H^t_{m_c+i,\star}\tilde{x}^t = 2 - \frac{1}{d_i}(\sum_t C^t_{i,\star}x^t) \leq 1$. Moreover, the optimality of $\{x^t\}_t$ guarantees that $\sum_t H^t_{0,\star}\tilde{x}^t = 2 - \frac{1}{\widehat{\mathsf{opt}}}(\sum_t \pi^t x^t) = 2 - \frac{\mathsf{opt}}{\widehat{\mathsf{opt}}} \leq 1$, where the last inequality follows from $\mathsf{opt} \geq \widehat{\mathsf{opt}}$. Thus, $\{\tilde{x}^t\}_t$ achieves load at most 1 for $\{H^t\}_t$.

(Lower bound.) Consider an optimal solution $\tilde{x}^1, \ldots, \tilde{x}^n \in \triangle^k$ for the load-balancing instance $\{H^t\}_t$, so $\lambda^* := \|\sum_t H^t x^t\|_{\max}$. Define $x^t = (\tilde{x}^t_1, \ldots, \tilde{x}^t_k)$ by dropping the last component from $\tilde{x}^t$.

For $i \in [m_p]$ we have $\lambda^* \geq \sum_t H^t_{i,\star}\tilde{x}^t = \frac{1}{b_i}\sum_t A^t_{i,\star}x^t$, and hence $\sum_t A^t\frac{x^t}{\lambda^*} \leq b$. Moreover, for $i \in [m_c]$ we have $\lambda^* \geq \sum_t H^t_{m_p+i,\star}\tilde{x}^t = 2 - \frac{1}{d_i}\sum_t C^t_{i,\star}x^t$, and hence $\sum_t C^t\frac{x^t}{\lambda^*} \geq \frac{2-\lambda^*}{\lambda^*}d \geq d$, where the last

inequality uses $\lambda^* \le 1$ which was proved in the first part of the claim. Thus, $\{\frac{x^t}{\lambda^*}\}_t$ is a feasible solution for the LP $\mathcal{L}$.

In addition, we have $\lambda^* \ge \sum_t H^t_{0,\star}\tilde{x}^t = 2 - \frac{1}{\widehat{\mathsf{opt}}}\sum_t \pi^t x^t$, and hence the solution $\{\frac{x^t}{\lambda^*}\}_t$ obtains value $\sum_t \pi^t \frac{x^t}{\lambda^*} \ge \frac{2-\lambda^*}{\lambda^*}\widehat{\mathsf{opt}}$. This must be at most the optimal value opt for the LP $\mathcal{L}$, so $\mathsf{opt} \ge \frac{2-\lambda^*}{\lambda^*}\widehat{\mathsf{opt}}$, or $\lambda^* \ge \frac{2\,\widehat{\mathsf{opt}}}{\widehat{\mathsf{opt}}+\mathsf{opt}} \ge \frac{1}{2}$, where the last inequality follows from $\frac{\mathsf{opt}}{2} \le \widehat{\mathsf{opt}} \le \mathsf{opt}$. This concludes the proof. □

Claim 4.3. *The instance $\{H^t\}_t$ is $(M,4)$-well-bounded for $M = \frac{2\varepsilon^2}{\log((m+1)/\delta)}$.*

Proof. Using the lower bound $\lambda^* \ge \frac{1}{2}$ from the previous claim, it suffices to show that each row of H^t contains entries that all lie in the interval $[-M, \frac{2}{n}]$, or all in the interval $[-\frac{2}{n}, M]$.

For the zeroth row, the generalized width lower bound implies $H^t_{0,j} \ge -\frac{\pi^t_j}{\widehat{\mathsf{opt}}} \ge -\frac{M\mathsf{opt}}{2\widehat{\mathsf{opt}}} \ge -M$, where the last inequality uses $\frac{\mathsf{opt}}{2} \le \widehat{\mathsf{opt}}$; since $H^t_{0,j} \le \frac{2}{n}$, it follows that $H^t_{0,j} \in [-M, \frac{2}{n}]$. For $i \in [m_p]$, we have $0 \le H^t_{i,j} \le \frac{a^t_{i,j}}{b_i} \le M$. And for $i \in [m_c]$, we have $H^t_{m_p+i,j} \le \frac{2}{n}$ and $H^t_{m_p+i,j} \ge -\frac{c^t_{i,j}}{d_i} \ge -M$. This concludes the proof. □

Proof of Theorem 4.1. Again let $M = \frac{2\varepsilon^2}{\log((m+1)/\delta)}$ and recall that $\varepsilon' = \sqrt{8}\varepsilon$. Since Claims 4.3 and 4.2 guarantee, respectively, that the instance $\{H^t\}_t$ is $(M,4)$-well-bounded and has optimal value $\lambda^* \ge \frac{3M\log((m+1)/\delta)}{(\varepsilon')^2}$, we can applying Theorem 3.3 to obtain that the solution $\{\tilde{\mathbf{x}}^t\}_t$ returned by the algorithm **expertLB** satisfies with probability at least $1-\delta$

$$\forall i = 0,\dots,m, \quad \sum_t \mathbf{H}^t_{i,\star}\tilde{\mathbf{x}}^t - \varepsilon'\sum_t |\mathbf{H}^t_{i,\star}\tilde{\mathbf{x}}^t| \le \lambda^*(1+4c_1\varepsilon') \le (1+4c_1\varepsilon'), \tag{4.11}$$

where the last inequality follows again from Claim 4.2. Let $\mathcal{E}$ be the event that this bound holds. For $i \in [m_p]$, since $H^t_{i,\star} = \frac{A^t_{i,\star}}{b_i}$ is non-negative, this means that conditioned on $\mathcal{E}$ we have $\sum_t \sum_{j\le k} \mathbf{A}^t_{ij}\tilde{\mathbf{x}}^t_j \le b_i\frac{(1+4c_1\varepsilon')}{(1-\varepsilon')}$ and so the solution $\{\widehat{\mathbf{x}}^t\}_t$ output by **LPviaLB** satisfies the packing constraints.

For $i = 0$ we have $|\mathbf{H}^t_{0,\star}\tilde{\mathbf{x}}^t| \le \frac{2}{n} + \frac{1}{\widehat{\mathsf{opt}}}\sum_{j\le k}\boldsymbol{\pi}^t_j\tilde{\mathbf{x}}^t_j$, and thus the left-hand side of inequality (4.11) is at least

$$\left(2 - \frac{1}{\widehat{\mathsf{opt}}}\sum_t\sum_{j\le k}\boldsymbol{\pi}^t_j\tilde{\mathbf{x}}^t_j\right) - \varepsilon'\left(2 + \frac{1}{\widehat{\mathsf{opt}}}\sum_t\sum_{j\le k}\boldsymbol{\pi}^t_j\tilde{\mathbf{x}}^t_j\right) = 2(1-\varepsilon') - \frac{(1+\varepsilon')}{\widehat{\mathsf{opt}}}\sum_t\sum_{j\le k}\boldsymbol{\pi}^t_j\tilde{\mathbf{x}}^t_j.$$

So conditioned on $\mathcal{E}$, inequality (4.11) gives

$$\sum_t\sum_{j\le k}\boldsymbol{\pi}^t_j\tilde{\mathbf{x}}^t_j \ge \widehat{\mathsf{opt}}\frac{(1-\varepsilon')(1-2\varepsilon'-c_1\varepsilon')}{(1+4c_1\varepsilon')(1+\varepsilon')} = (1-O(\varepsilon'))\,\widehat{\mathsf{opt}},$$

where the last inequality uses the fact that ε (and hence ε') is at most a sufficiently small constant. Therefore, conditioned on $\mathcal{E}$, we have that the solution output by **LPviaLB** satisfies $\sum_t \boldsymbol{\pi}^t\widehat{\mathbf{x}}^t = \frac{(1-\varepsilon')}{(1+4c_1\varepsilon')}\sum_t\sum_{j\le k}\boldsymbol{\pi}^t_j\tilde{\mathbf{x}}^t_j \ge (1-O(\varepsilon'))\,\widehat{\mathsf{opt}}$.

An identical proof shows that conditioned on $\mathcal{E}$ each covering constraint is $O(\varepsilon')$-approximately satisfied, i.e., $\sum_t \mathbf{C}^t\widehat{\mathbf{x}}^t \ge (1-O(\varepsilon'))\,d$. This concludes the proof of the theorem. □

5 Estimating the optimal value based on sampled LPs

In ths preceding section, we saw how to solve a PCMC LP $\mathcal{L}$ given an estimate of $\mathsf{opt}(\mathcal{L})$. Now we show how to obtain such estimates via "sampled LPs". For that, we require the LP to satisfy a *stability* condition. Loosely speaking this asks that the covering constraints are not "very tight" — if we reduce the covering requirements by a little, the optimal value should not increase by a lot. First, a definition to simplify notation:

Definition 5.1. *Given a PCMC LP $\mathcal{L}$ and $\varepsilon \in [0,1)$, we define $\mathcal{L}(1-\varepsilon)$ as the LP obtained by multiplying the right-hand side of the covering constraints by $(1-\varepsilon)$, but keeping the packing constraints unchanged.*

Notice that the feasible solutions for $\mathcal{L}(1-\varepsilon)$ are exactly the ε-feasible solutions for $\mathcal{L}$. Now we formally define the stability property.

Definition 5.2 (Stability)**.** *A PCMC LP $\mathcal{L}$ is called* (ε,σ)-stable *if any optimal ε-feasible solution for it has value at most $(1+\sigma\varepsilon)\cdot\mathsf{opt}(\mathcal{L})$, i.e., $\mathsf{opt}(\mathcal{L}(1-\varepsilon)) \le (1+\sigma\varepsilon)\cdot\mathsf{opt}(\mathcal{L})$.*

The following lemma gives a natural sufficient condition for an LP to be stable.

Lemma 5.3. *If a PCMC LP $\mathcal{L}$ is such that $\mathcal{L}(1+\frac{1}{\sigma})$ is feasible, then $\mathcal{L}$ is (ε,σ)-stable for every $\varepsilon \in (0,1]$.*

Proof. Consider any $\varepsilon \ge 0$. Let $\{\bar{x}^t\}_t$ be a feasible solution for $\mathcal{L}(1+\frac{1}{\sigma})$ and let $\{\underline{x}^t\}_t$ be an optimal solution for $\mathcal{L}(1-\varepsilon)$. Define $\tilde{x}^t = \frac{\varepsilon}{\varepsilon+1/\sigma}\bar{x}^t + (1-\frac{\varepsilon}{\varepsilon+1/\sigma})\underline{x}^t$. It is easy to verify that $\{\tilde{x}^t\}_t$ is feasible for $\mathcal{L}$ and has value

$$\sum_{t=1}^{n} \pi^t \tilde{x}^t \ge \left(1 - \frac{\varepsilon}{\varepsilon+1/\sigma}\right)\mathsf{opt}(\mathcal{L}(1-\varepsilon)),$$

where the inequality uses the non-negativity of the π^t's. This implies that $\mathsf{opt}(\mathcal{L}) \ge (1-\frac{\varepsilon}{\varepsilon+1/\sigma})\mathsf{opt}(\mathcal{L}(1-\varepsilon))$, which is equivalent to $\mathsf{opt}(\mathcal{L}(1-\varepsilon)) \le (1+\varepsilon\sigma)\mathsf{opt}(\mathcal{L})$. This concludes the proof. □

The idea behind our estimator for $\mathsf{opt}(\mathcal{L})$ in the random-order model has been used in prior works on **packing-only** LPs, e.g., [DH09, FHK+10, AWY14, MR12]: observe the first s random blocks $(\boldsymbol{\pi}^1,\mathbf{A}^1,\mathbf{C}^1),\ldots,(\boldsymbol{\pi}^s,\mathbf{A}^s,\mathbf{C}^s)$ of the LP, and form the *sampled LP* with these blocks and right-hand side scaled by a factor of $\frac{s}{n}$:

$$\begin{aligned} \max\ & \textstyle\sum_{t=1}^{s} \boldsymbol{\pi}^t\mathbf{x}^t & \\ st\ & \textstyle\sum_{t=1}^{s} \mathbf{A}^t\mathbf{x}^t \le \frac{s}{n}\cdot b & \\ & \textstyle\sum_{t=1}^{s} \mathbf{C}^t\mathbf{x}^t \ge \frac{s}{n}\cdot d & \\ & \mathbf{x}^t \in \blacktriangle^k \quad \forall t \in [s]. & \end{aligned} \tag{5.12}$$

By the random-order assumption, the blocks $(\boldsymbol{\pi}^t,\mathbf{A}^t,\mathbf{C}^t)$ are an $\frac{s}{n}$-sample of the underlying LP $\mathcal{L}$, and hence should be "representative" of the latter. Consequently, computing an optimal (ε-feasible) solution for this sampled LP should give a good estimate of $\frac{s}{n}\cdot\mathsf{opt}(\mathcal{L})$.

In this work, we extend this idea beyond packing-only LPs, and allow for covering constraints as well. This introduces complications, but the stability property defined above allows us to address them. Let us define sampled LPs in the generality that we need.

Definition 5.4 (Restricted LP)**.** *Given a PCMC LP $\mathcal{L}$, and a subset I of $[n]$, the* restricted LP $\mathcal{L}^I$ *is obtained by retaining only the columns of $\mathcal{L}$ that belong to I, and setting the right-hand side* $\mathsf{rhs}(\mathcal{L}^I)$ *to be* $\frac{|I|}{n}\mathsf{rhs}(\mathcal{L})$*, namely*

$$\begin{aligned} \max \ & \textstyle\sum_{t\in I} \pi^t x^t \\ st \ & \textstyle\sum_{t\in I} A^t x^t \le \frac{|I|}{n}\cdot b \\ & \textstyle\sum_{t\in I} C^t x^t \ge \frac{|I|}{n}\cdot d \\ & x^t \in \blacktriangle^k \quad \forall t \in I. \end{aligned} \tag{$\mathcal{L}^I$}$$

Notice that the restricted LP $\boldsymbol{\mathcal{L}}^{\{1,\dots,s\}}$ is exactly the sampled LP of (5.12). Moreover, since the sequence of random variables $(\boldsymbol{\pi}^1, \mathbf{A}^1, \mathbf{C}^1), \dots, (\boldsymbol{\pi}^n, \mathbf{A}^n, \mathbf{C}^n)$ is exchangeable (i.e., every subsequence of same size has the same joint distribution), $\boldsymbol{\mathcal{L}}^I$ has the same distribution as the sampled LP $\boldsymbol{\mathcal{L}}^{\{1,\dots,s\}}$ for every s-subset $I \subseteq [n]$ (only the indices of the variables are different).

In the following, we will need a more refined control over the ratio of item profits to opt, we define the *width* of a PCMC LP $\mathcal{L}$ as

$$\min\left\{ \min_{t,i,j} \frac{b_i}{a^t_{ij}},\ \min_{t,i,j} \frac{d_i}{c^t_{ij}} \right\} \ . \tag{5.13}$$

Note that unlike the generalized width defined in (4.10), the width does not depend on $\frac{\pi^t_{i,j}}{\mathsf{opt}}$. Finally, for a vector $v \in \mathbb{R}^n$ and a subset $I \subseteq [n]$, we use $v|_I = (v_i)_{i\in I}$ to denote the vector which keeps only the coordinates in I.

The following is the main result of this section: The first two parts say that with good probability a solution to the original LP translates to a solution to the sampled LP with value not much lower than the expected value. This is via a straightforward application of Bernstein's inequality. The latter two parts show that the optimal value for the sampled LP is not much higher either, both in expectation and with high probability. These parts use (ε, σ)-stability to control the variance of the (dual of the) sampled LP. Since the proof is slightly technical we defer it to Appendix C.

Lemma 5.5 (Sampled LP Lemma)**.** *For $\varepsilon, \delta \in (0,1)$ and $\rho \ge 1$, consider a PCMC LP $\mathcal{L}$ with width at least $\frac{32\log(m/\delta)}{\varepsilon^2}$ and such that all item values $\pi^t_{i,j}$ are at most $\rho \cdot \mathsf{opt}(\mathcal{L})(\frac{\varepsilon^2}{32\log(m/\delta)})$. Let $s \ge \varepsilon^2 n$ and let I be a s-subset of $[n]$. Then for $\varepsilon' = \varepsilon\sqrt{\frac{n}{s}}$ we have the following guarantees for the random LP $\boldsymbol{\mathcal{L}}$:*

(a) If $\bar{\mathbf{x}}$ is feasible for $\boldsymbol{\mathcal{L}}$, then $(1-\frac{\varepsilon'}{2})\cdot \bar{\mathbf{x}}|_I$ is feasible for $\boldsymbol{\mathcal{L}}^I(1-\varepsilon')$ with probability at least $1-2\delta$.

(b) With probability at least $1-4\delta$, $\mathsf{opt}(\boldsymbol{\mathcal{L}}^I(1-\varepsilon')) \ge (1-\frac{\rho\varepsilon'}{2})\cdot\frac{s}{n}\,\mathsf{opt}(\mathcal{L})$.

Moreover, if $\mathcal{L}$ is (ε,σ)-stable, then the following upper bounds hold:

(c) With probability at least $1-2\delta$, $\mathsf{opt}(\boldsymbol{\mathcal{L}}^I(1-\varepsilon')) \le (1+2\sigma\varepsilon'+\rho\varepsilon')\cdot\frac{s}{n}\,\mathsf{opt}(\mathcal{L})$

(d) $\mathbb{E}\big[\mathsf{opt}(\boldsymbol{\mathcal{L}}^I(1-\varepsilon'))\big] \le (1+\sigma\varepsilon')\cdot\frac{s}{n}\,\mathsf{opt}(\mathcal{L})$.

6 Solving packing/covering LPs with unknown OPT

We can now combine the algorithm from §4 that needs estimates on opt with the algorithm from §5 that estimates of opt to get almost optimal ε-feasible online solutions to packing-covering multiple-choice LPs. We are going to use the same *doubling* strategy introduced in [AWY14] (and used in [MR12]) in the packing-only case. We give the details since covering constraints introduce technicalities that preclude applying previous results directly. To simplify notation we use the shorthand $\boldsymbol{\mathcal{L}}^{\le k}$ to denote the sampled LP $\boldsymbol{\mathcal{L}}^{\{1,\dots,k\}}$, and similarly $\boldsymbol{\mathcal{L}}^{>k} := \boldsymbol{\mathcal{L}}^{\{k+1,\dots,n\}}$.

6.1 One-time learning OPT

The *one-time learning* (**OTL**) algorithm below uses the first half of the instance to estimate the value of opt and then runs the algorithm **LPviaLB** on the second half.

Algorithm 3 . One-time learning (**OTL**) algorithm

Input: Reals $\varepsilon, \delta, \sigma > 0$, number n of items, and right-hand sides b and d.

(Estimation phase) Let $\varepsilon' = \varepsilon\sqrt{2}$. For time steps $1, \ldots, \frac{n}{2}$ observe blocks $(\boldsymbol{\pi}^1, \mathbf{A}^1, \mathbf{C}^1)$, $\ldots, (\boldsymbol{\pi}^{n/2}, \mathbf{A}^{n/2}, \mathbf{C}^{n/2})$ and compute the estimate $\widehat{\mathsf{opt}} := \frac{1-2\varepsilon'}{1+3\sigma} \cdot \mathsf{opt}(\mathcal{L}^{\leq \frac{n}{2}}(1-\varepsilon'))$. Output $\bar{\mathbf{x}}^t = 0$ for all these time steps.

(Optimization phase) For the remaining time steps $\frac{n}{2}+1, \ldots, n$, run algorithm **LPviaLB** with the computed estimate $\widehat{\mathsf{opt}}$ over the scaled remaining LP $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')$, outputting the solution $\bar{\mathbf{x}}^{n/2+1}, \ldots, \bar{\mathbf{x}}^n$ returned by **LPviaLB**.

The following theorem shows that the solution computed is almost feasible and almost optimal for the second half $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ of the instance.

Theorem 6.1. *Consider a feasible PCMC LP $\mathcal{L}$. Consider $\varepsilon > 0$ at most a sufficiently small constant, and $\delta \in (0, \varepsilon]$, and assume that the generalized width of $\mathcal{L}$ is at least $\frac{32\log(m/\delta)}{\varepsilon^2}$. Also assume that $\mathcal{L}$ is (ε, σ)-stable for $\sigma \in [1, \frac{1}{20\sqrt{2}\varepsilon}]$. Then with probability at least $1 - 13\delta$, the algorithm **OTL** returns an online solution $\{\bar{\mathbf{x}}^t\}_{t=n/2}^n$ that is $(c_4\varepsilon)$-feasible for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ and has value*

$$\sum_{t=1}^{n} \boldsymbol{\pi}^t \bar{\mathbf{x}}^t \geq (1 - c_5\sigma\varepsilon)\, \mathsf{opt}\big(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}\big(1-\sqrt{2}\varepsilon\big)\big)$$

for constants c_4, c_5.

The proof is conceptually direct: it uses Lemma 5.5 to assert the quality of the estimate $\widehat{\mathsf{opt}}$ defined, and then employs Theorem 4.1 to obtain guarantees on the computed solution $\{\bar{\mathbf{x}}^t\}_{t=n/2+1}^n$. The details are in Appendix D.

Observe that this theorem does not give an almost feasible solution for the entire LP $\boldsymbol{\mathcal{L}}$, only for $\boldsymbol{\mathcal{L}}^{>n/2}$. Also it only achieves value $\approx \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>n/2}) \approx \frac{1}{2}\mathsf{opt}(\mathcal{L})$. To do better we use the doubling strategy to dynamically learn opt.

6.2 Dynamically learning OPT

Our final algorithm for PCMC LPs is called *dynamic learning algorithm* (**DLA**). The idea of this algorithm is to use the blocks seen in the first εn time steps to obtain an estimate of $\mathsf{opt}(\mathcal{L})$ and feed it to **LPviaLB** to compute solutions for time steps $\varepsilon n + 1, \ldots, \varepsilon n$. Then it uses the blocks seen up to this time (i.e. the first $2\varepsilon n$ blocks) to obtain a better estimate of $\mathsf{opt}(\mathcal{L})$ and again feeds it to **LPviaLB**, obtaining solutions for time steps $2\varepsilon n+1, \ldots, 4\varepsilon n$; the algorithm proceeds in this doubling fashion until the end. Equivalently, we can define it as repeatedly using algorithm **OTL** from the previous section:

Algorithm 4 . Dynamic learning algorithm (**DLA**)

Input: Reals $\varepsilon, \delta, \sigma > 0$, number n of items, and right-hand sides b and d.

For $i = 0, \ldots, \log \varepsilon^{-1}$, let $S_i := \{1, \ldots, 2^i \varepsilon n\}$ and let $\varepsilon_i = \varepsilon\sqrt{n/|S_i|} = \sqrt{\varepsilon/2^i}$

for $i = 1, \ldots, \log \varepsilon^{-1}$ **do**

(Phase i) Run the algorithm **OTL** with parameters $\varepsilon_i, \delta, \sigma$ over the instance $\boldsymbol{\mathcal{L}}^{S_i}$, namely the first $2^i \varepsilon n$ times steps of instance $\boldsymbol{\mathcal{L}}$, and let $\left(0, \ldots, 0, \bar{\mathbf{x}}^{|S_{i-1}|+1}, \bar{\mathbf{x}}^{|S_{i-1}|+2}, \ldots, \bar{\mathbf{x}}^{|S_i|}\right)$ be the returned solution

end for

Output $\left(0, \ldots, 0, \bar{\mathbf{x}}^{|S_0|+1}, \bar{\mathbf{x}}^{|S_0|+2}, \ldots, \bar{\mathbf{x}}^{n}\right)$, namely the "union" of the solutions computed above

Algorithm **DLA** finally delivered the promised guarantees for PCMC LPs: for stable LPs with not too small generalized width, it obtains an $O(\varepsilon)$-feasible $(1 - O(\varepsilon))$-optimal solution with high probability. (This corresponds to Theorem 1.2 in the introduction.)

Theorem 6.2. *Consider a feasible PCMC LP $\mathcal{L}$. Consider $\varepsilon > 0$ at most a sufficiently small constant ε_0, and $\delta \in (0, \varepsilon]$, and assume that the generalized width of $\mathcal{L}$ is at least $\frac{32 \log(m/\delta)}{\varepsilon^2}$. Also assume that $\mathcal{L}$ is (ε, σ)-stable for $\sigma \in [1, \frac{1}{20\sqrt{2}\varepsilon}]$. Then with probability at least $1 - 23\delta \log(1/\varepsilon)$, the algorithm **DLA** returns an online solution $\{\bar{\mathbf{x}}^t\}_t$ that is $(c_6\varepsilon)$-feasible for $\boldsymbol{\mathcal{L}}$ and has value*

$$\sum_{t=1}^{n} \boldsymbol{\pi}^t \bar{\mathbf{x}}^t \geq (1 - c_7 \sigma \varepsilon)\, \mathsf{opt}(\mathcal{L})$$

for constants c_6, c_7.

The main element for proving this theorem is to obtain guarantees for the solutions returned by **OTL** in each phase i of the algorithm **DLA**. This is done by showing that with good probability the random instance $\boldsymbol{\mathcal{L}}^{S_i}(1-\varepsilon_i)$ is $(\varepsilon_i, O(\sigma))$-stable (using the bounds from Lemma 5.5) so we can employ Theorem 6.1 to obtain guarantees for **OTL**; a full proof is presented in Appendix E.

Lemma 6.3. *Let $\bar{\mathbf{x}}(i) = \left(0, \ldots, 0, \bar{\mathbf{x}}^{|S_{i-1}|+1}, \bar{\mathbf{x}}^{|S_{i-1}|+2}, \ldots, \bar{\mathbf{x}}^{|S_i|}\right)$ be the partial solution computed in phase i of algorithm **DLA**. Then with probability at least $1 - 19\delta$, $\bar{x}(i)$ is $((c_4 + 1)\varepsilon_i)$-feasible for $\boldsymbol{\mathcal{L}}^{S_i \setminus S_{i-1}}$ and has value at least $\sum_{t=1}^{|S_i|} \boldsymbol{\pi}^t \bar{\mathbf{x}}(i)^t \geq (1 - 7c_5\sigma\varepsilon_i)\, \mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i \setminus S_{i-1}}(1 - \sqrt{2}\varepsilon_i))$.*

We now use this lemma on every phase of **DLA** to obtain the proof of Theorem 6.2.

Proof of Theorem 6.2. Let $\mathcal{E}$ be the event that for all $i = 1, \ldots, \log(\frac{1}{\varepsilon})$, the bound from Lemma 6.3 holds and

$$\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i \setminus S_{i-1}}(1 - \sqrt{2}\varepsilon_i)) \geq \left(1 - \varepsilon_i\right) \cdot \frac{|S_i \setminus S_{i-1}|}{n}\, \mathsf{opt}(\mathcal{L}) = \left(1 - \varepsilon_i\right) \varepsilon\, 2^{i-1}\, \mathsf{opt}(\mathcal{L})$$

also holds. Using Lemmas 5.5 and 6.3 and taking a union bound over $i = 1, \ldots, \log(\frac{1}{\varepsilon})$, we get that event $\mathcal{E}$ holds with probability at least $1 - 23\delta \log(\frac{1}{\varepsilon})$ (indeed we can apply Lemma 5.5 since $\varepsilon_i \geq \varepsilon$ implies that the generalized width of $\boldsymbol{\mathcal{L}}^{S_i \setminus S_{i-1}}$ is at least $\frac{32 \log(m/\delta)}{\varepsilon^2} \geq \frac{32 \log(m/\delta)}{\varepsilon_i^2}$, and $\varepsilon_i^2 \leq \varepsilon$ implies $|S_i \setminus S_{i-1}| \geq \varepsilon_i^2 n$). In addition, notice that the definition of the ε_i's gives

$$\sum_{i=1}^{\log(1/\varepsilon)} \varepsilon_i 2^{i-1} = \frac{\sqrt{\varepsilon}}{2} \sum_{i=1}^{\log(1/\varepsilon)} \sqrt{2^i} \leq \frac{\sqrt{\varepsilon}}{2} \cdot \frac{\sqrt{2}^{\log(1/\varepsilon)+1}}{\sqrt{2} - 1} = \frac{1}{\sqrt{2}\,(\sqrt{2} - 1)} \leq 2. \tag{6.14}$$

Then using the guarantees from Lemma 6.3, it is easy to see that under the event $\mathcal{E}$ the solution $\{\bar{\mathbf{x}}^t\}_t$ output by **DLA** satisfies the covering constraints within a factor of $(1-\varepsilon(1+2c_4))$:

$$\sum_{t=1}^{n} \mathbf{C}^t \bar{\mathbf{x}}^t \geq d \cdot \sum_{i=1}^{\log(1/\varepsilon)} (1-(c_4+1)\varepsilon_i) \frac{|S_i \setminus S_{i-1}|}{n} = d \cdot \left(1-\varepsilon-\varepsilon(c_4+1)\sum_{i=1}^{\log(1/\varepsilon)} \varepsilon_i 2^{i-1}\right) \geq d \cdot (1-\varepsilon(3+2c_4)).$$

Similar, and even simpler, calculations show that the solution $\{\bar{\mathbf{x}}^t\}_t$ also satisfies the packing constraints under $\mathcal{E}$, and thus is $(3+2c_4)$-feasible for $\boldsymbol{\mathcal{L}}$. Finally, since under $\mathcal{E}$ we have that $\sum_{t=|S_{i-1}|+1}^{|S_i|} \bar{\mathbf{x}}^t \geq (1-7c_5\sigma\varepsilon_i)(1-\varepsilon_i)\,\varepsilon 2^{i-1}\,\mathsf{opt}(\mathcal{L})$, adding over all phases i gives

$$\begin{aligned}\sum_{t=1}^{n} \boldsymbol{\pi}^t \bar{\mathbf{x}}^t &\geq \varepsilon\,\mathsf{opt}(\mathcal{L}) \sum_{i=1}^{\log(1/\varepsilon)} (1-7c_5\sigma\varepsilon_i)\big(1-\varepsilon_i\big)\,2^{i-1} \geq \varepsilon\,\mathsf{opt}(\mathcal{L}) \sum_{i=1}^{\log(1/\varepsilon)} \Big(1-\frac{c_7}{2}\sigma\varepsilon_i\Big)\,2^{i-1} \\ &\geq \mathsf{opt}(\mathcal{L})\,(1-c_7\varepsilon\sigma)\end{aligned}$$

for some constant c_7, where in the second inequality we also used the fact $\sigma \geq 1$. This concludes the proof of the theorem. □

7 Solving packing LPs with no assumption on item values

Finally, we turn to removing all assumptions about the magnitudes of the profits π_t. To achieve this, we consider packing-only LPs (PLPs) of the following form:

$$\begin{aligned} \max\ & \textstyle\sum_{t=1}^{n} \pi_t x_t \\ st\ & Ax \leq b \\ & x \in [0,1]^n, \end{aligned} \tag{PLP}$$

for non-negative π, A, and b. Observe the absence of covering constraints, and also of multiple-choice constraints—each load matrix A^t is just a column vector, so at each step one more column of the matrix A is revealed to us. In this setting we now find a solution of value $(1-O(\varepsilon))\,\mathsf{opt}$, without any assumptions about the magnitudes of the values π_t.

The presence of items with large value π_t compared to opt rules out using algorithms **DLA** or **LPviaLB** directly, since we may violate the required generalized width bounds. We would like to show a threshold value such that items having values higher than the threshold can be directly added to our solution (i.e., we can set $x_t = 1$ for these items) such that: (i) the total value of these items plus the LP on the remaining items is still a $(1-\varepsilon)$-approximation to the original optimal value (i.e., we do not lose much value by adding these items); (ii) the individual values of items in the remaining LP are $O\big(\frac{\varepsilon^2}{\log(m/\delta)}\big) \cdot \mathsf{opt}$, and (iii) the right-hand side of the remaining LP is still large. Properties (ii) and (iii) would guarantee that the generalized width of the remaining LP is $\Omega\big(\frac{\log(m/\delta)}{\varepsilon^2}\big)$ so we can employ the algorithm **LPviaLB**, for example, on the remaining LP.

Indeed, our algorithm proceeds essentially using this approach. The difference is that we can only show a threshold such that the remaining items have values smaller than $O\big(\frac{\varepsilon^2}{\log(m/\delta)}\big)$ times the *original LP value*, not the new LP value. Thus, having taken high-value items into our solution, the value of the new "skimmed-off" LP may have fallen drastically, so we may get no bounds on the generalized width. Yet, the value of the remaining LP being small implies the already-picked items have a lot of value, which means even a poor approximation on the remaining LP suffices. This approach introduces additional complications that are handled in the remainder of §7.

Let us establish some notation. Denote the set of PLPs with n columns by PLP_n, and the right-hand side vector b of a PLP $\mathcal{L}$ by $\mathsf{rhs}(\mathcal{L})$. Also, recall that the *width* (defined in (5.13)) of a packing LP is

$\max_{i,j}\{\frac{b_i}{a_{ij}}\}$. Define $B := \frac{\log((m+1)/\delta)}{\varepsilon^2}$, and will want the width of all LPs to be $\Omega(B)$. *By scaling, we will assume that entries in each column A^t lie in $[0,1]$ and that the right-hand sides b_i's are at least the width of the LP.*

7.1 Estimating the high-value items

The first piece of the puzzle is to show the existence of a good skimming threshold τ such that blindly adding to the solution the items with value more than τ satisfies the properties (i)–(iii) above. We first formally define the "remaining LP" mentioned above.

Definition 7.1 (Skimmed LP)**.** *For $\mathcal{L} \in \mathit{PLP}_n$ and a threshold $\tau \geq 0$, consider the high-value items $H = \{t : \pi_t > \tau\}$. The* skimmed LP *$\mathit{skim}_\tau(\mathcal{L})$ is then obtained from $\mathcal{L}$ by setting to zero the value of all items in H and setting the right-hand side to be $\mathsf{rhs}(\mathit{skim}_\tau(\mathcal{L})) := \mathsf{rhs}(\mathcal{L}) - \sum_{t\in H} A^t$. (I.e., this is essentially the residual LP if we hard-code $x_t = 1$ for all t with $\pi_t > \tau$.)*

The following gives a formal definition of a good skimming threshold (it is useful to think of $\Delta = \Theta(\mathsf{opt}(\mathcal{L}))$, which corresponds to the discussion above).

Definition 7.2 (Good Skimming Threshold)**.** *A value $\tau \in \mathbb{R}$ is a* Δ-good skimming threshold *for $\mathcal{L} \in \mathit{PLP}_n$ if:*

(S1) $\mathsf{opt}(\mathit{skim}_\tau(\mathcal{L})) + \sum_{t:\pi_t>\tau} \pi_t \geq (1-\varepsilon)\,\mathsf{opt}(\mathcal{L})$,
(S2) $\mathsf{rhs}(\mathit{skim}_\tau(\mathcal{L})) \geq \frac{1}{2}\mathsf{rhs}(\mathcal{L})$, *and*
(S3) $\tau \leq \frac{24\cdot\Delta}{B}$.

The starting point is the following simple but crucial claim, which gives a "perfect" skimming threshold τ — one where we do not lose *any value* by blindly picking items of higher value. It is a generalization of the "value/weight" rule for solving single knapsack problems, and says that every optimal fractional solution must pick all the sufficiently high valued items.

Lemma 7.3 (Perfect Threshold)**.** *For a feasible PLP $\mathcal{L}$ and for every $\tau \geq \frac{\mathsf{opt}(\mathcal{L})}{\min_i b_i}$, we have*

$$\mathsf{opt}(\mathit{skim}_\tau(\mathcal{L})) + \sum_{t:\pi_t>\tau} \pi_t = \mathsf{opt}(\mathcal{L}).$$

Proof. The dual of $\mathcal{L}$ is given by

$$\mathcal{D} = \min\left\{ \sum_{i\in[m]} b_i p_i + \sum_{t\in[n]} y_t \;\middle|\; \langle p, A^t\rangle + y_t \geq \pi_t \;\; \forall t \in [n],\; p, y \geq 0 \right\},$$

where the variable y_t is relative to the primal constraint $x_t \leq 1$.

Let $(x^*, (p^*, y^*))$ be an optimal primal-dual pair for $\mathcal{L}$ and $\mathcal{D}$. We claim that x^* selects all items with $\pi_t > \|p^*\|_1$. To see this, observe that $\pi_t > \langle p^*, A^t\rangle$ means $y_t^* > 0$ and hence $x_t^* = 1$ by complementary slackness. But since the entries of A^t are upper bounded by 1, we have $\langle p^*, A^t\rangle \leq \|p^*\|_1$. Thus, $\pi_t > \|p^*\|_1 \geq \langle p^*, A^t\rangle$ implies $x_t^* = 1$, giving the claim.

This means that

$$\mathsf{opt}\big(\mathsf{skim}_{\|p^*\|_1}(\mathcal{L})\big) + \sum_{t:\pi_t>\|p^*\|_1} \pi_t = \mathsf{opt}(\mathcal{L}).$$

Now it suffices to show that $\|p^*\|_1 \leq \frac{\mathsf{opt}(\mathcal{L})}{\min_i b_i}$. This is easy from strong duality: $\mathsf{opt}(\mathcal{L}) = \mathsf{opt}(\mathcal{D}) \geq \sum_{i\in[m]} b_i p_i^* \geq (\min_i b_i)\cdot\|p^*\|_1$, which completes the proof. ∎

Lemma ensures $\frac{\mathsf{opt}(\mathcal{L})}{\min_i b_i}$ satisfies criteria (S1) and (S3) for being a $\Theta(\mathsf{opt}(\mathcal{L}))$-good skimming threshold notice. But if we assume that $\mathcal{L}$ has width at least B, the slightly higher threshold $\tau = \frac{2\mathsf{opt}(\mathcal{L})}{\min_i b_i} \leq \frac{2\mathsf{opt}(\mathcal{L})}{B}$ satisfies (S2) as well. Indeed, at most $\frac{B}{2}$ items with value more than τ (else picking $\frac{B}{2}+1$ of them would satisfy the constraints of the LP $\mathcal{L}$ and give value strictly greater than $\mathsf{opt}(\mathcal{L})$), and thus $\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})) \geq \mathsf{rhs}(\mathcal{L}) - \frac{B}{2}\mathbb{1} \geq \frac{1}{2}\mathsf{rhs}(\mathcal{L})$. Therefore, this threshold τ is indeed a "perfect" $\Theta(\mathsf{opt}(\mathcal{L}))$-good skimming threshold for $\mathcal{L}$.

Hence we want to estimate τ in order to pick up the high-value items in an online fashion, i.e., we seem to need to estimate $\mathsf{opt}(\mathcal{L})$ with good probability (in the presence of unboundedly large item values π_t). And hence we seem to be back at square one. However, recall that this is only if we want a perfect threshold. By excluding a small number of items with the largest values π_t, we now show how to find a good-enough skimming threshold. To make this precise, for $K \in \mathbb{N}$ let $\mathcal{L}_{<K}$ denote the LP obtained by zeroing out the values π_t of the K highest valued items (without changing the right hand side).[5]

First we show that the optimal value of this modified LP $\mathcal{L}_{<K}$ can be (over)estimated by using the optimal value of the sampled LP $\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}$.

Lemma 7.4. *For $\mathcal{L} \in \mathit{PLP}_n$ with width at least $16B$ and $\varepsilon \leq \frac{1}{32}$, and $K = 128\varepsilon B$,*

$$\Pr\left[\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}) \geq \frac{\mathsf{opt}(\mathcal{L}_{<K})}{8}\right] \geq 1 - 6\delta.$$

Proof. Let M be the value of the K^{th} largest value of π. We use Lemma 5.5(b) on $\mathcal{L}_{<K}$ (using the fact that all item values in $\mathcal{L}_{<K}$ are at most M and that since there are no covering constraints $\mathcal{L}_{<K}(1-\varepsilon') = \mathcal{L}_{<K}$) to get with probability at least $1-4\delta$

$$\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}) \geq \mathsf{opt}((\boldsymbol{\mathcal{L}}_{<K})^{\leq \frac{n}{4}}) \geq \frac{\mathsf{opt}(\mathcal{L}_{<K})}{4} - 8MB\varepsilon = \frac{\mathsf{opt}(\mathcal{L}_{<K})}{4} - \frac{KM}{16}. \tag{7.15}$$

Moreover, we claim that $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}) \geq (KM)/16$ with probability at least $1-2\delta$. For that, let $\mathbf{T}$ be the indices of the top $K = 128\varepsilon B$ values of $\boldsymbol{\pi}$ and consider the solution $\widehat{\mathbf{x}}$ that sets for all $t \in [n/4]$, $\widehat{\mathbf{x}}_t = 1$ if $t \in \mathbf{T}$ and $\widehat{\mathbf{x}}_t = 0$ otherwise. We know that $\mathsf{rhs}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}) = \frac{1}{4}\mathsf{rhs}(\mathcal{L}) \geq 4B\mathbb{1}$; moreover, this solution $\widehat{\mathbf{x}}$ can (at most) pick all $K = 128\varepsilon B \leq 4B$ elements of $\mathbf{T}$ (since $\varepsilon \leq 1/32$), so $\widehat{\mathbf{x}}$ is feasible for $\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}$. The expected number of items picked by $\widehat{\mathbf{x}}$ is $\mathbb{E}[|\mathbf{T} \cap [n/4]|] = K/4$, and the Bernstein's inequality Corollary 2.4 implies $\Pr(|\mathbf{T} \cap [n/4]| < K/16) \leq 2\exp(-\varepsilon B) \leq 2\frac{\delta}{m}$. Each item with $\widehat{\mathbf{x}}_t = 1$ has value at least M, so with probability at least $1-2\delta$

$$\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}) \geq \sum_{t \in \mathbf{T} \cap [n/4]} \boldsymbol{\pi}_t \widehat{\mathbf{x}}_t \geq \frac{KM}{16}. \tag{7.16}$$

Taking a union bound we have that with probability at least $1-6\delta$ inequalities (7.15) and (7.16) hold simultaneously. When this holds we can take the average of these inequalities, and so with probability at least $1-6\delta$ $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}}) \geq \frac{\mathsf{opt}(\mathcal{L}_{<K})}{8}$, which concludes the proof. □

Next, the optimal value of the modified LP $\mathcal{L}_{<K}$ can be used to obtain a good skimming threshold.

Lemma 7.5. *Consider a feasible PLP $\mathcal{L}$ with width at least B. Suppose $\varepsilon \leq \frac{1}{3}$ and K is at most $\frac{\varepsilon}{2}$ times the width of $\mathcal{L}$. Then any $\tau \geq \frac{3\mathsf{opt}(\mathcal{L}_{<K})}{B}$ is a $\frac{\tau B}{24}$-good skimming threshold for $\mathcal{L}$.*

[5] We assume items have distinct values for simplicity; one can break ties consistently and get the same result.

Proof. We first show that condition (1) of a good skimming threshold holds, namely $\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L})) + \sum_{t:\pi_t>\tau} \pi_t \geq (1-\varepsilon)\,\mathsf{opt}(\mathcal{L})$. Let $H = \{t : \pi_t > \tau\}$ be the high-value items, and let $T \subseteq [n]$ denote the index set of the K items with largest value π_t (the "top-value" items).

Since the width of $\mathcal{L}_{<K}$ is the same as that of $\mathcal{L}$ (which is at least B), we have that $\tau \geq \frac{3\mathsf{opt}(\mathcal{L}_{<K})}{\min_i \mathsf{rhs}(\mathcal{L}_{<K})_i}$ and thus from Lemma 7.3 we have

$$\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}_{<K})) + \sum_{t\notin T, t\in H} \pi_t = \mathsf{opt}(\mathcal{L}_{<k}).$$

Adding all the top valued items T to both sides gives

$$\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}_{<K})) + \sum_{t\in H} \pi_t + \sum_{t\in T, t\notin H} \pi_t = \mathsf{opt}(\mathcal{L}_{<k}) + \sum_{t\in T} \pi_t \geq \mathsf{opt}(\mathcal{L}). \tag{7.17}$$

To connect the left-hand side of this expression with $\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}))$ we claim that

$$\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L})) \geq (1-\varepsilon)\,\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}_{<K})) + \sum_{t\in T, t\notin H} \pi_t. \tag{7.18}$$

To see this, consider an optimal solution x^* for $\mathsf{skim}_\tau(\mathcal{L}_{<K})$ (which only accrues value from items in $\overline{H}$), and construct the vector $\widehat{x}$ by setting $\widehat{x}_t = (1-\varepsilon)x^*_t$ for $t \in \overline{H}\cap\overline{T}$, $\widehat{x}_t = 1$ for $t \in \overline{H}\cap T$, and $\widehat{x}_t = 0$ otherwise. Notice that $\widehat{x}$ is feasible for $\mathsf{skim}_\tau(\mathcal{L})$: First, $\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})) \geq \mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L}_{<K})) - K\mathbb{1}$, since only the K items in T can be additionally skimmed in $\mathcal{L}$ compared to $\mathcal{L}_{<k}$. Thus, the usage of the solution $\widehat{x}$ is

$$\begin{aligned}\sum_{t:t\leq\tau} A^t\widehat{x}_t &\leq (1-\varepsilon)\sum_{t:t\leq\tau} A^t x^*_t + K\mathbb{1} \leq (1-\varepsilon)\,\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L}_{<K})) + K\mathbb{1}\\ &\leq (1-\varepsilon)\,\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})) + 2K\mathbb{1} \leq \mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})),\end{aligned}$$

where the inequality $2K\mathbb{1} \leq \varepsilon\,\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L}))$ used in the last step follows from the fact that $\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})) \geq \frac{1}{2}\mathsf{rhs}(\mathcal{L})$ (condition (2) proved below) and $\mathsf{rhs}(\mathcal{L}) \geq \frac{2K}{\varepsilon}\mathbb{1}$ (assumption). Therefore, the solution $\widehat{x}$ is feasible for $\mathsf{skim}_\tau(\mathcal{L})$. Since it has value equal to the right-hand side of inequality (7.18), it proves the validity of this inequality.

Putting inequalities (7.17) and 7.18 together and using the non-negativity of the π_t's gives $\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}))+\sum_{t:\pi_t>\tau}\pi_t \geq (1-\varepsilon)\mathsf{opt}(\mathcal{L})$, and thus condition (1) holds.

Condition (3) in the definition of good skimming threshold trivially holds, so to conclude the proof we show condition (2), namely $\mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})) \geq \frac{1}{2}\mathsf{rhs}(\mathcal{L})$. As argued before, there are at most $\frac{B}{3}$ items in $H\cap\overline{T}$, otherwise we could pick $\frac{B}{3}+1$ items in this set to obtain a feasible solution for $\mathcal{L}_{<K}$ (since $A^t \leq \mathbb{1}$ for all t) with value more than $\mathsf{opt}(\mathcal{L}_{<K})$, a contradiction. Therefore, the right-hand side of $\mathsf{skim}_\tau(\mathcal{L})$ equals

$$\mathsf{rhs}(\mathcal{L}) - \sum_{t\in H} A^t \geq \mathsf{rhs}(\mathcal{L}) - \left(|H\cap\bar{T}| + |T|\right)\cdot\mathbb{1} \geq \mathsf{rhs}(\mathcal{L}) - \left(\frac{B}{3}+K\right)\cdot\mathbb{1}.$$

Since the width of $\mathcal{L}$ is at least B and at least $\frac{2K}{\varepsilon}$, this inequality gives $\mathsf{skim}_\tau(\mathcal{L}) \geq \mathsf{rhs}(\mathcal{L})(1-\frac{1}{3}-\frac{\varepsilon}{2}) \geq \frac{1}{2}\mathsf{rhs}(\mathcal{L})$, where the last inequality uses $\varepsilon \leq \frac{1}{3}$. This concludes the proof of the lemma. □

Therefore, putting Lemmas 7.4 and 7.5 together, we get that with probability at least $1-6\delta$ the value $\boldsymbol{\tau} := \frac{24}{B}\,\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq\frac{n}{4}})$ is an $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\mathcal{L}$; in fact, as will be more useful, it is an $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\boldsymbol{\mathcal{L}}^{>\frac{n}{4}}$ and $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ (this simply uses the additional fact that $\mathsf{opt}(\mathcal{L}_{<K})$ is at least $\mathsf{opt}((\boldsymbol{\mathcal{L}}^{>\frac{n}{4}})_{<K})$ and $\mathsf{opt}((\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})_{<K})$).

Lemma 7.6 (Finding Good Skimming Threshold)**.** *Consider $\mathcal{L} \in \mathit{PLP}_n$ with width at least $512B$, $\varepsilon \leq \frac{1}{32}$, and $\delta > 0$. Define $\boldsymbol{\tau} := \frac{24}{B}\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq\frac{n}{4}})$. Then $\boldsymbol{\tau}$ is an $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\boldsymbol{\mathcal{L}}^{>\frac{n}{4}}$ and $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ with probability at least $1-6\delta$.*

7.2 Estimating opt and right-hand side of a skimmed LP

In order to solve the skimmed LP $\mathsf{skim}_\tau(\mathcal{L})$ using the algorithm **LPviaLB**, we need to obtain estimates for the optimal value of this LP, as well as its right-hand side (which is now random). Crucially, since $\mathsf{skim}_\tau(\mathcal{L})$ has only "small-value" items (and we assume $\mathcal{L}$ has large width), these quantities can be estimated by looking at that sampled LP $\boldsymbol{\mathcal{L}}^{\leq \frac{n}{3}}$ just like we did in §5 for packing-covering LPs; a full proof is deferred to Appendix F.

Lemma 7.7. *Let $\varepsilon > 0$ be at most a sufficiently small constant. Given $\mathcal{L} \in \mathsf{PLP}_n$ with width at least $40B$, let τ be an $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\mathcal{L}$. Then with probability at least $1 - 4\delta$,*

$$\mathsf{rhs}\big(\mathit{skim}_\tau(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{3}})\big) \in (1 \pm 3\varepsilon)\,\frac{1}{2}\,\mathsf{rhs}\big(\mathit{skim}_\tau(\boldsymbol{\mathcal{L}}^{> \frac{n}{3}})\big)$$

and with probability at least $1 - 16\delta$,

$$\mathsf{opt}\big(\mathit{skim}_\tau(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{3}})\big) \in (1 \pm 3\varepsilon)\,\frac{1}{2}\,\mathsf{opt}\big(\mathit{skim}_\tau(\boldsymbol{\mathcal{L}}^{> \frac{n}{2}})\big) \pm 33\varepsilon B\tau.$$

Next, we show we can replace this fixed good skimming threshold τ by the estimated one $\boldsymbol{\tau}$ given by Lemma 7.6. For that, we condition that the set of items that appear on times $1, \ldots, \frac{n}{4}$ is such that $\boldsymbol{\tau}$ is a $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\boldsymbol{\mathcal{L}}^{> \frac{n}{4}}$, and then apply Lemma 7.7 to the LP $\boldsymbol{\mathcal{L}}^{> \frac{n}{4}}$; taking an average over all such conditionings then gives the following.

Corollary 7.8. *Let $\varepsilon > 0$ be at most a sufficiently small constant, and let $\mathcal{L} \in \mathsf{PLP}_n$ have width at least $512B$. Define $\boldsymbol{\tau} := \frac{24}{B}\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq \frac{n}{4}})$ as in Lemma 7.6. Then with probability at least $1 - 14\delta$ all of the following hold simultaneously:*

$$\boldsymbol{\tau} \text{ is a } \mathsf{opt}(\mathcal{L})\text{-good skimming threshold for } \boldsymbol{\mathcal{L}}^{> \frac{n}{2}}$$

$$\mathsf{rhs}\big(\mathit{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{(\frac{n}{4}, \frac{n}{2}]})\big) \in (1 \pm 3\varepsilon)\,\frac{1}{2}\,\mathsf{rhs}\big(\mathit{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{> \frac{n}{2}})\big)$$

$$\mathsf{opt}\big(\mathit{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{(\frac{n}{4}, \frac{n}{2}]})\big) \in (1 \pm 3\varepsilon)\,\frac{1}{2}\,\mathsf{opt}\big(\mathit{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{> \frac{n}{2}})\big) \pm 33\varepsilon B\boldsymbol{\tau}.$$

Finally, since we only have an **estimate** of the right-hand side of $\mathsf{skim}_\tau(\mathcal{L})$, we need to ensure that the algorithm **LPviaLB** can handle this. Thankfully this is the case: The following corollary of Theorem 4.1 follows directly from the fact that if we reduce the right-hand side of a PLP by a factor of $(1 - 6\varepsilon)$, the optimal value reduces by at most a factor of $(1 - 6\varepsilon)$ (since we can scale an original optimal solution by $(1 - 6\varepsilon)$ and obtain at least this much value).

Corollary 7.9 (of Theorem 4.1)**.** *Consider a feasible PLP $\mathcal{L}$. Consider $\varepsilon > 0$ at most a sufficiently small constant ε_0, and $\delta \in (0, \varepsilon]$, and assume that the generalized width of $\mathcal{L}$ is at least $\frac{2\log((1+m)/\delta)}{\varepsilon^2}$. Suppose the algorithm **LPviaLB** is given an estimate $\widehat{\mathsf{opt}} \in [\frac{\mathsf{opt}}{2}, \mathsf{opt}]$ and an approximation $\widehat{b} \in [(1 - 6\varepsilon)b, b]$ to the packing right-hand side. Then **LPviaLB** finds an online solution $\{\mathbf{x}^t\}_t$ that with probability at least $1 - \delta$ is feasible for $\boldsymbol{\mathcal{L}}$ and has value $\sum_t \boldsymbol{\pi}^t \mathbf{x}^t \geq (1 - c_3\varepsilon - 6\varepsilon)\,\widehat{\mathsf{opt}}$.*

7.3 Modified one-time learning algorithm

Now we present the *modified one-time learning* (**mOTL**) algorithm, which at a high-level does the natural thing: 1) uses the first $\frac{n}{2}$ time steps (i.e. $\boldsymbol{\mathcal{L}}^{\leq \frac{n}{2}}$) to estimate the required parameters (good skimming threshold $\boldsymbol{\tau}$, optimal value and right-hand side of $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{> \frac{n}{2}})$); 2) in time steps $> \frac{n}{2}$, picks all high-value items $\boldsymbol{\pi}_t > \boldsymbol{\tau}$ and feeds the remaining LP $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{\frac{n}{2}})$ to algorithm **LPviaLB**. Here is

a formal description of the algorithm **mOTL**, which has one small additional feature to simplify its analysis: we truncate the solution obtained for $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$ to ensure it is always feasible for estimated right-hand side $\widehat{\mathbf{b}}$.

Algorithm 5 . Modified one-time learning (**mOTL**) algorithm

Input: Reals $\varepsilon, \delta, \sigma > 0$, number n of items, and right-hand side b.

(Estimation) For time steps $1, \ldots, \frac{n}{2}$ observe blocks $(\boldsymbol{\pi}^1, \mathbf{A}^1, \mathbf{C}^1), \ldots, (\boldsymbol{\pi}^{n/2}, \mathbf{A}^{n/2}, \mathbf{C}^{n/2})$ and compute the estimates

$$\begin{aligned}
\boldsymbol{\tau} &:= \frac{24}{B}\,\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\leq\frac{n}{4}}) \\
\widehat{\mathbf{O}} &:= \frac{2}{1+3\varepsilon}\left[\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{(\frac{n}{4},\frac{n}{2}]})) - 33\varepsilon B\boldsymbol{\tau}\right] && \text{(estimate of } \mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))\text{)} \\
\widehat{\mathbf{b}} &:= \frac{2}{1+3\varepsilon}\,\mathsf{rhs}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{(\frac{n}{4},\frac{n}{2}]})) && \text{(estimate of } \mathsf{rhs}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))\text{)}
\end{aligned}$$

(Approximate solution for $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$) Define $\varepsilon' = \varepsilon \max\left\{1, \sqrt{2B\boldsymbol{\tau}/\widehat{\mathbf{O}}}\right\}$. Run algorithm **LPviaLB** with the computed estimates $\widehat{\mathbf{O}}$ and $\widehat{\mathbf{b}}$, and parameters ε', δ, over the skimmed remaining LP $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$, letting $\tilde{\mathbf{x}}_{\frac{n}{2}+1}, \ldots, \tilde{\mathbf{x}}_n$ be the returned solution. Let $\mathbf{t}_{\text{last}}$ be the largest index such that $\sum_{t=\frac{n}{2}+1}^{\mathbf{t}_{\text{last}}} \mathbf{A}^t(1-6\varepsilon)\tilde{\mathbf{x}}_t \leq \widehat{\mathbf{b}}$.

(Output above solution + high-value items, truncated) Let $\mathbf{H} = \{t > \frac{n}{2} : \boldsymbol{\pi}_t > \boldsymbol{\tau}\}$ be the high-value items within the time periods $\frac{n}{2}+1, \ldots, n$. Output the solution $\bar{\mathbf{x}}_{\frac{n}{2}+1}, \ldots, \bar{\mathbf{x}}_n$ defined by

$$\bar{\mathbf{x}}_t = \begin{cases} 1 & \text{if } t \in \mathbf{H} \\ (1-6\varepsilon)\,\tilde{\mathbf{x}}_t & \text{if } t \notin \mathbf{H}, t \in \left[\frac{n}{2}+1, \mathbf{t}_{\text{last}}\right] \\ 0 & \text{if } t \notin \mathbf{H}, t > \mathbf{t}_{\text{last}}. \end{cases}$$

The next theorem states the guarantees that we have for algorithm **mOTL**.

Theorem 7.10. *Consider $\varepsilon > 0$ smaller than a sufficiently small constant, and $\delta \in (0, \varepsilon]$. Consider a feasible LP $\mathcal{L} \in \mathsf{PLP}_n$ with width at least $512B$. Then with probability at least $1 - 15\delta$, the algorithm **mOTL** returns a solution $\bar{\mathbf{x}}_{\frac{n}{2}+1}, \ldots, \bar{\mathbf{x}}_n$ that is feasible for the LP $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ and has value at least*

$$\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}) - c_6\varepsilon\,\mathsf{opt}(\mathcal{L})$$

for some constant c_6.

Here is some intuition for the proof. Recall that the main point of the skimming operation is to try to control the *generalized* width of the skimmed LP $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$; let us see what we have accomplished in this direction. Suppose (via Corollary 7.8) that the estimate $\boldsymbol{\tau}$ computed is an $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$. Property (2) of a good skimming threshold implies that the *width* of $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$ is the desired $\Omega(B) = \Omega\big(\frac{\log(m/\delta)}{\varepsilon^2}\big)$ (if we assume that the width of $\mathcal{L}$ is $\Omega(B)$). However, for the item values, notice that the ratio of $\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))$ divided by the largest item value in $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$ is at least $\frac{\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))}{\boldsymbol{\tau}} = \Omega(B)\,\frac{\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))}{\mathsf{opt}(\mathcal{L})} = \Omega\big(\frac{\log(m/\delta)}{\varepsilon^2}\big)\frac{\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))}{\mathsf{opt}(\mathcal{L})}$. If the skimmed LP has large optimal value (namely $\Omega(\mathsf{opt}(\mathcal{L}))$), then the skimmed LP has generalized width $\Omega\big(\frac{\log(m/\delta)}{\varepsilon^2}\big)$ and we can guarantee that **LPviaLB** returns an $(1 - O(\varepsilon))$-approximation for it; since $\boldsymbol{\tau}$ is a good skimming

threshold, this $(1-O(\varepsilon))$-approximation together with the high-value items guarantee that algorithm **mOTL** obtains value $\approx \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$.

However, if the skimmed LP has a small optimal value, the generalized width is *not* within the desired $\Omega\big(\frac{\log(m/\delta)}{\varepsilon^2}\big)$ range and we lose the guarantee for algorithm **LPviaLB**. But the generalized width of $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$ is always in the range $\Omega\big(\frac{\log(m/\delta)}{(\varepsilon'')^2}\big)$ for $\varepsilon'' = \Theta(\varepsilon)\sqrt{\frac{\mathsf{opt}(\mathcal{L})}{\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))}}$, and thus we can get a $(1-O(\varepsilon''))$-approximation (notice the similarity of ε'' and the ε' used in the algorithm). Moreover, when $\varepsilon'' \gg \varepsilon$ this means that the high-value items of $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ have a large contribution to the optimal value $\mathsf{opt}(\mathcal{L})$; since these items are picked by our algorithm, our solution should again have large overall value.

Proof of Theorem 7.10. Let σ be a realization of $\boldsymbol{\mathcal{L}}^{\leq\frac{n}{2}}$ where the guarantees from Corollary 7.8 hold (the probability of obtaining such realization is at least $1-14\delta$). In this scenario, the solution $\bar{\mathbf{x}}$ output by **mOTL** is feasible for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$: we have $\widehat{\mathbf{b}} \leq \mathsf{rhs}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))$ and thus the choice of stopping time $\mathbf{t}_{\text{last}}$ makes $((1-6\varepsilon)\tilde{\mathbf{x}}_{\frac{n}{2}+1}, \dots, (1-6\varepsilon)\tilde{\mathbf{x}}_{\mathbf{t}_{\text{last}}}, 0, \dots, 0)$ feasible for $\mathsf{rhs}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))$; the definition of skimmed LP then guarantees that the output solution $\bar{\mathbf{x}}$ (which consists of these $\tilde{\mathbf{x}}_t$'s plus high-value items) fits the budgets of $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$, thus we get feasibility.

As for the value val of the solution under scenario σ, we have two different cases depending on the value of $\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))$; let C be a sufficiently large constant ($C = \frac{132\times48\times2}{1-9\varepsilon}$ suffices).

Case 1: $\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) < \varepsilon C\mathsf{opt}(\mathcal{L})$. Since $\boldsymbol{\tau}$ is a $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$, the high-value items already give value

$$\sum_{t\in\mathbf{H}} \boldsymbol{\pi}_t \geq (1-\varepsilon)\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}) - \mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) \geq \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}) - \varepsilon(C+1)\mathsf{opt}(\mathcal{L}),$$

and thus $\mathsf{val} \geq \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}) - \varepsilon(C+1)\mathsf{opt}(\mathcal{L})$.

Case 2: $\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) \geq \varepsilon C\mathsf{opt}(\mathcal{L})$. We claim that the requirements of Corollary 7.9 are fulfilled (with approximation parameter ε'), so we can apply it to obtain guarantees for the solution $\tilde{\mathbf{x}}$ for $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$; most conditions can be easily checked, we argue about the less trivial ones. First, since Corollary 7.8 holds for scenario σ, $\widehat{\mathbf{O}} \leq \mathsf{opt}\big(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})\big)$; then $\varepsilon' \geq \varepsilon\sqrt{2B\boldsymbol{\tau}/\mathsf{opt}\big(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})\big)}$ and reorganizing this expression gives that the LP $\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})$ has generalized width at least $\frac{2\log((m+1)/\delta)}{(\varepsilon')^2}$. Second, the estimate $\widehat{\mathbf{O}}$ is at least $\frac{\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))}{2}$:

$$\widehat{\mathbf{O}} \geq \frac{1-3\varepsilon}{1+3\varepsilon}\,\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) - O(\varepsilon B\boldsymbol{\tau}) \geq \frac{1-3\varepsilon}{1+3\varepsilon}\,\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) - O(\varepsilon\mathsf{opt}(\mathcal{L})) \geq \frac{1}{2}\,\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})),$$

where the last inequality uses the assumption that $\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) \geq \varepsilon C\mathsf{opt}(\mathcal{L})$ for a sufficiently large C and that ε is at most a sufficiently small constant. Finally, using this fact plus the bound $\boldsymbol{\tau} \leq \frac{24\mathsf{opt}(\mathcal{L})}{B}$, we have $\varepsilon' \leq \varepsilon_0$:

$$\varepsilon' = \varepsilon\max\left\{1, \sqrt{2B\boldsymbol{\tau}/\widehat{\mathbf{O}}}\right\} \leq \varepsilon\max\left\{1, \sqrt{\frac{4\times24\mathsf{opt}(\mathcal{L})}{\mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}))}}\right\} \leq \max\big\{\varepsilon, O(\sqrt{\varepsilon})\big\} \leq \varepsilon_0,$$

where the second inequality follows from our assumption on Case 2, and the last inequality follows from the fact ε is a sufficiently small constant.

Thus, we can apply Corollary 7.9 to $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ and get that, conditioned on σ, with probability at least $1-\delta$ the solution $\tilde{\mathbf{x}}_{\frac{n}{2}+1}, \ldots, \tilde{\mathbf{x}}_n$ is feasible for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$ and has value

$$\sum_{t>\frac{n}{2}:t\notin\mathbf{H}} \boldsymbol{\pi}_t \tilde{\mathbf{x}}_t \geq (1 - c_3\varepsilon' - 6\varepsilon')\,\widehat{\mathbf{O}} \geq \widehat{\mathbf{O}} - \varepsilon(c_3+6)\max\left\{\widehat{\mathbf{O}}, \sqrt{2B\boldsymbol{\tau}\widehat{\mathbf{O}}}\right\} \geq \widehat{\mathbf{O}} - O(\varepsilon)\,\mathsf{opt}(\mathcal{L})$$
$$\geq \mathsf{opt}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) - O(\varepsilon)\,\mathsf{opt}(\mathcal{L}),$$

where the second inequality because $\widehat{\mathbf{O}} \leq \mathsf{opt}(\mathcal{L})$ (a consequence of an upper bound on $\widehat{\mathbf{O}}$ proved above) and the fact $\boldsymbol{\tau}$ is $\mathsf{opt}(\mathcal{L})$-good skimming threshold, and the last inequality uses the above upper bound on $\widehat{\mathbf{O}}$. Also, the feasibility of this solution also implies (together with the conditioning on σ) that

$$\sum_{t>\frac{n}{2}} \mathbf{A}^t \tilde{\mathbf{x}}_t \leq \mathsf{rhs}(\mathsf{skim}_{\boldsymbol{\tau}}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}})) \leq \frac{1}{1-6\varepsilon}\,\widehat{\mathbf{b}},$$

and hence the stopping time $\mathbf{t}_{\text{last}}$ is set to be equal to n. Thus, the value val obtained by the algorithm uses the full solution $\tilde{\mathbf{x}}_{\frac{n}{2}+1}, \ldots, \tilde{\mathbf{x}}_n$, and again using the fact $\boldsymbol{\tau}$ is a good skimming threshold we get

$$\mathsf{val} = \sum_{t>\frac{n}{2}:t\notin\mathbf{H}} \boldsymbol{\pi}_t \tilde{x}_t + \sum_{t\in\mathbf{H}} \boldsymbol{\pi}_t \geq \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}) - O(\varepsilon)\,\mathsf{opt}(\mathcal{L}).$$

Since this all happened with probability at least $1-\delta$ conditioned on a scenario σ where Corollary 7.8 holds, taking an expectation over all such scenarios gives that the above guarantee for val holds with overall probability at least $1-15\delta$. This concludes the proof of the theorem. □

7.4 Final algorithm: Modified Dynamic Learning

The final algorithm to solve packing LPs, *modified dynamic learning* (**mDLA**), is exactly as the dynamic learning algorithm from §6.2 with the only change that it uses the modified One-time learning algorithm **mOTL** instead of **OTL**.

The following theorem shows that this algorithm **mDLA** has the desired guarantees.

Theorem 7.11. *Consider $\varepsilon > 0$ smaller than a sufficiently small constant, and $\delta \in (0, \varepsilon]$. Consider a feasible LP $\mathcal{L} \in \mathsf{PLP}_n$ with width at least $1024\,\frac{\log((m+1)/\delta)}{\varepsilon^2}$. Then there is an event $\mathcal{E}$ that holds with probability at least $1 - c_7\,\delta\log\frac{1}{\varepsilon}$ such that under $\mathcal{E}$ the solution $\{\bar{\mathbf{x}}_t\}_t$ returned by algorithm **mDLA** is feasible for the LP $\boldsymbol{\mathcal{L}}$ and has expected value conditioned on $\mathcal{E}$ at least $(1 - c_8\delta\log\varepsilon^{-1})\,\mathsf{opt}(\mathcal{L})$, for some constants c_7, c_8. In particular, the truncated solution $\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2, \ldots, \bar{\mathbf{x}}_{\mathbf{t}'_{last}}, 0, \ldots, 0$, where $\mathbf{t}'_{last}$ is the largest index that makes this solution feasible for $\boldsymbol{\mathcal{L}}$, has expected value at least $(1 - c_9\delta\log e^{-1})\,\mathsf{opt}(\mathcal{L})$ for some constant c_9.*

The proof of this theorem is similar to the proof of the guarantee for the algorithm **DLA** and the changes required —which arise from the fact that we now do not have good bounds on the ratios $\pi_t/\mathsf{opt}(\mathcal{L})$— are indicated in Appendix G.

References

[ABFP13] Yossi Azar, Umang Bhaskar, Lisa Fleischer, and Debmalya Panigrahi. Online mixed packing and covering. In *SODA*, pages 85–100, 2013.

[ABH11] Jacob Abernethy, Peter L. Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *COLT 2011 - The 24th Annual Conference on Learning Theory, June 9-11, 2011, Budapest, Hungary*, pages 27–46, 2011. URL: `http://www.jmlr.org/proceedings/papers/v19/abernethy11b/abernethy11b.pdf`.

[AD14] Shipra Agrawal and Nikhil R. Devanur. Bandits with concave rewards and convex knapsacks. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 989–1006, New York, NY, USA, 2014. ACM. URL: http://doi.acm.org/10.1145/2600057.2602844, doi:10.1145/2600057.2602844.

[AD15] Shipra Agrawal and Nikhil R. Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1405–1424, 2015. URL: http://dx.doi.org/10.1137/1.9781611973730.93, doi:10.1137/1.9781611973730.93.

[AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012. URL: http://www.theoryofcomputing.org/articles/v008a006, doi:10.4086/toc.2012.v008a006.

[AWY14] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *Operations Research*, 62(4):876–890, 2014. URL: http://dx.doi.org/10.1287/opre.2014.1289, arXiv:http://dx.doi.org/10.1287/opre.2014.1289, doi:10.1287/opre.2014.1289.

[BIK07] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA '07*, pages 434–443, 2007.

[BIKK07] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *APPROX-RANDOM*, pages 16–28, 2007.

[BIKK08] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exchanges*, 7(2), 2008.

[BN09] Niv Buchbinder and Joseph Naor. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends in Theoretical Computer Science*, 3(2-3):93–263, 2009.

[CBL06] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[CHW12] Kenneth L. Clarkson, Elad Hazan, and David P. Woodruff. Sublinear optimization for machine learning. *J. ACM*, 59(5):23:1–23:49, 2012.

[DF80] P. Diaconis and D. Freedman. Finite exchangeable sequences. *The Annals of Probability*, 8(4):745–764, 08 1980. URL: http://dx.doi.org/10.1214/aop/1176994663.

[DH09] Nikhil R. Devenur and Thomas P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In John Chuang, Lance Fortnow, and Pearl Pu, editors, *ACM Conference on Electronic Commerce*, pages 71–78. ACM, 2009.

[DJSW11] Nikhil R. Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In Yoav Shoham, Yan Chen, and Tim Roughgarden, editors, *ACM Conference on Electronic Commerce*, pages 29–38. ACM, 2011.

[Dyn63] E. B. Dynkin. The optimum choice of the instant for stopping a Markov process. *Soviet Math. Dokl*, 4, 1963.

[ESF14] Reza Eghbali, Jon Swenson, and Maryam Fazel. Exponentiated subgradient algorithm for online optimization under the random permutation model. *CoRR*, abs/1410.7171, 2014. URL: http://arxiv.org/abs/1410.7171.

[FHK+10] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Clifford Stein. Online stochastic packing applied to display ad allocation. In Mark de Berg and Ulrich Meyer, editors, *ESA (1)*, volume 6346 of *Lecture Notes in Computer Science*, pages 182–194. Springer, 2010.

[Fre75] David A. Freedman. On tail probabilities for martingales. *Annals of Probability*, 3:100–118, 1975. doi:10.1214/aop/1176996452.

[GK07] Naveen Garg and Jochen Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *SIAM J. Comput.*, 37(2):630–652 (electronic), 2007. URL: http://dx.doi.org/10.1137/S0097539704446232.

[GM08] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adWords. In Shang-Hua Teng, editor, *SODA*, pages 982–991. SIAM, 2008.

[Haz06] Elad Hazan. Approximate convex optimization by online game playing. *CoRR*, abs/cs/0610119, 2006. URL: http://arxiv.org/abs/cs/0610119.

[Kha04] Rohit Khandekar. *Lagrangian Relaxation Based Algorithms for Convex Programming Problems*. PhD thesis, Indian Institute of Technology Delhi, March 2004. URL: http://people.cse.iitd.ac.in/~rohitk/research/thesis.ps.gz.

[Kle05] Robert Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '05, pages 630–631, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics. URL: http://portal.acm.org/citation.cfm?id=1070432.1070519.

[KRTV13] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. *CoRR*, abs/1311.2578, 2013. To appear in STOC 2014. URL: http://arxiv.org/abs/1311.2578.

[KV05] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291 – 307, 2005. URL: http://www.sciencedirect.com/science/article/pii/S0022000004001394, doi:http://dx.doi.org/10.1016/j.jcss.2004.10.016.

[MR12] Marco Molinaro and R. Ravi. Geometry of online packing linear programs. In *Automata, languages, and programming. Part I*, volume 7391 of *Lecture Notes in Comput. Sci.*, pages 701–713. Springer, Heidelberg, 2012. URL: http://dx.doi.org/10.1007/978-3-642-31594-7_59.

[MSVV07] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. AdWords and generalized online matching. *J. ACM*, 54(5):Art. 22, 19, 2007. URL: http://dx.doi.org/10.1145/1284320.1284321.

[Pru98] Alexander R. Pruss. A maximal inequality for partial sums of finite exchangeable sequences of random variables. *Proceedings of the American Mathematical Society*, 126(6):1811–1819, 1998.

[PST95] Serge A. Plotkin, David B. Shmoys, and Éva Tardos. Fast approximation algorithms for fractional packing and covering problems. *Math. Oper. Res.*, 20(2):257–301, 1995. URL: http://dx.doi.org/10.1287/moor.20.2.257.

[Sim89] R. W. Simpson. Using network flow techniques to find shadow prices for market and seat inventory control. Technical report, 1989. MIT Flight Transportation Laboratory Memorandum M89-1.

[SS11] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011. URL: http://dx.doi.org/10.1561/2200000018, doi:10.1561/2200000018.

[TR98] Kalyan Talluri and Garrett Van Ryzin. An analysis of bid-price controls for network revenue management. *Manage. Sci.*, 44(11):1577–1593, November 1998. URL: http://dx.doi.org/10.1287/mnsc.44.11.1577.

[vdVW96] Aad W. van der Vaart and Jon A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics. URL: http://link.springer.com/book/10.1007/978-1-4757-2545-2.

[Wil92] Elizabeth Louise Williamson. *Airline network seat inventory control: methodologies and revenue impacts*. PhD thesis, M.I.T., 1992.

[You01] Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In *42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)*, pages 538–546. IEEE Computer Soc., Los Alamitos, CA, 2001.

A Maximal Bernstein for Sampling Without Replacement

We state the Levy-type maximal inequality for exchangeable random variables obtained in [Pru98].

Theorem A.1 (Theorem 1 of [Pru98], with $\gamma = 2$)**.** *Let $Y_1, Y_2, \ldots, Y_n$ be an exchangeable sequence of random variables taking values in a Banach space. Then for every $\lambda \geq 0$ and every $k \leq n/2$, we have*

$$\Pr\left(\max_{j\leq k}\left\|\sum_{i\leq j} Y_i\right\| > \lambda\right) \leq 15\Pr\left(\left\|\sum_{i\leq k} Y_i\right\| > \frac{\lambda}{24}\right).$$

Proof of Lemma 2.5. The sequence $(\sum_{i\leq j} X_i - j\mu)_{i=1}^n$ is exchangeable, so applying the theorem above and then Theorem 2.3 we have

$$\Pr\left(\max_{j\leq k}\left|\sum_{i\leq j} X_i - j\mu\right| > \alpha\right) \leq 15\Pr\left(\left|\sum_{i\leq k} X_i - k\mu\right| > \frac{\alpha}{24}\right) \leq 30\exp\left(-\frac{(\alpha/24)^2}{2k\sigma^2 + (\alpha/24)}\right).$$

□

B Useful Properties of (ε, σ)-Stability

For any ε' and subset $I \subseteq [n]$, the dual of $\mathcal{L}^I(1-\varepsilon')$ is

$$\mathcal{D}^I(1-\varepsilon') := \left\{\begin{array}{ll} \min & \frac{k}{n}\langle\alpha, b\rangle - (1-\varepsilon')\frac{k}{n}\langle\beta, d\rangle + \sum_{t\in I}\gamma_t \\ & \left\langle A^t_{\star,j}, \alpha\right\rangle - \left\langle C^t_{\star,j}, \beta\right\rangle + \gamma_t \geq \pi^t_j \quad \forall j, \quad \forall t \in I \\ & \alpha, \beta, \gamma \geq 0 \end{array}\right. \tag{B.19}$$

where again we use $A^t_{\star,j}$ to denote the jth column of A^t, and similarly for C^t.

The next essentially follows from the fact that $\eta \mapsto \mathsf{opt}(\mathcal{L}(1-\eta))$ is concave, and that its derivative is obtained by looking at an optimal dual solution.

Lemma B.1. *Let $\mathcal{L}$ be an (ε, σ)-stable packing/covering LP. Then:*

1. *For every $\varepsilon' \geq \varepsilon$, $\mathcal{L}$ is (ε', σ)-stable*
2. *If $(\tilde\alpha, \tilde\beta, \tilde\gamma)$ is an optimal solution for the dual of $\mathcal{L}(1-\varepsilon)$, then $\left\langle\bar\beta, d\right\rangle \leq \sigma \cdot \mathsf{opt}(\mathcal{L})$.*

Proof. Take an optimal solution $(\tilde\alpha, \tilde\beta, \tilde\gamma)$ for $\mathcal{D}(1-\varepsilon)$. Notice that for every η, $(\bar\alpha, \bar\beta, \bar\gamma)$ is feasible for $\mathcal{D}(1-\varepsilon+\eta)$ and hence $\mathsf{opt}(\mathcal{D}(1-\varepsilon+\eta)) \leq \langle\bar\alpha, b\rangle - (1-\varepsilon+\eta)\left\langle\bar\beta, d\right\rangle + \sum_t \bar\gamma = \mathsf{opt}(\mathcal{D}(1-\varepsilon)) - \eta\left\langle\bar\beta, d\right\rangle$. Using Strong Duality, this gives $\mathsf{opt}(\mathcal{L}(1-\varepsilon)) \geq \mathsf{opt}(\mathcal{L}(1-\varepsilon+\eta)) + \eta\left\langle\bar\beta, d\right\rangle$.

Setting $\eta = \varepsilon$, (ε, σ)-stability gives that $\left\langle\bar\beta, d\right\rangle \leq \sigma \cdot \mathsf{opt}(\mathcal{L})$, proving Part 2. Also, setting $\eta = \varepsilon - \varepsilon'$ we get $\mathsf{opt}(\mathcal{L}(1-\varepsilon')) \leq \mathsf{opt}(\mathcal{L}(1-\varepsilon)) + (\varepsilon' - \varepsilon)\sigma \cdot \mathsf{opt}(\mathcal{L})$; again using (ε, σ)-stability of $\mathcal{L}$ we get $\mathsf{opt}(\mathcal{L}(1-\varepsilon')) \leq (1+\varepsilon'\sigma)\mathsf{opt}(\mathcal{L})$ for $\varepsilon' \geq \varepsilon$, concluding the proof of Part 1. □

C Proof of Lemma 5.5

Since the LP has width at least $\frac{32\log(m/\delta)}{\varepsilon^2}$, by scaling the constraints we assume without loss of generality that its right-hand sides are all exactly equal to $\vartheta := \frac{32\log(m/\delta)}{\varepsilon^2}$ and the matrices A^t, C^t on the left-hand side have entries in $[0, 1]$.

Part (a). Consider a feasible solution $\bar{\mathbf{x}}$ for $\boldsymbol{\mathcal{L}}$. Since for all i, $\sum_t \mathbf{A}_i^t \bar{\mathbf{x}}^t \le \vartheta$, from Bernstein's inequality (Corollary 2.4) we have

$$\Pr\left(\sum_{t\in I} \mathbf{A}_i^t \bar{\mathbf{x}}^t \ge \frac{s}{n}\vartheta + \tau\vartheta\right) \le 2\exp\left(-\min\left\{\frac{\tau^2\vartheta^2}{8s\vartheta/n}, \frac{\tau\vartheta}{2}\right\}\right) \le 2\exp\left(-\frac{\log(m/\delta)}{\varepsilon^2}\min\left\{\frac{4\tau^2 n}{s}, 16\tau\right\}\right);$$

setting $\tau = (\varepsilon/2)\sqrt{s/n}$ makes the right-hand side at most $2\delta/m$ (this uses the fact $\varepsilon \le 1$ and $s \ge \varepsilon^2 n$). Then taking a union bound over all $i \in [m_p]$, we have $\sum_{t\in I} \mathbf{A}^t \bar{\mathbf{x}}^t \le (1 + \frac{\varepsilon'}{2})\frac{s}{n}\vartheta\mathbb{1}$ with probability at least $1 - 2\delta(m_p/m)$.

Similarly, fix $i \in \{m_p + 1, \dots, m\}$ and let $\mu = \mathbb{E}[\sum_{t\in I} \mathbf{C}_i^t \bar{\mathbf{x}}^t]$; then using the fact $\mu \ge (s/n)\vartheta$,

$$\Pr\left(\sum_{t\in I} \mathbf{C}_i^t \bar{\mathbf{x}}^t \le \mu - \tau\mu\frac{n}{s}\right) \le 2\exp\left(-\min\left\{\frac{\tau^2\mu(n/s)^2}{8}, \frac{\tau\mu(n/s)}{2}\right\}\right) \le 2\exp\left(-\frac{\log(m/\delta)}{\varepsilon^2}\min\left\{\frac{4\tau^2 n}{s}, 16\tau\right\}\right).$$

Setting $\tau = (\varepsilon/2)\sqrt{s/n}$ and taking a union bound over all such i's, we get $\sum_{t\in I} \mathbf{C}^t \bar{\mathbf{x}}^t \ge (1 - \frac{\varepsilon'}{2})\frac{s}{n}\vartheta\mathbb{1}$ with probability at least $1 - 2\delta(m_c/m)$.

Then, using the fact that $(1 - \frac{\varepsilon'}{2})(1 + \frac{\varepsilon'}{2}) \le 1$, with probability at least $1 - 2\delta$ we have that $(1 - \frac{\varepsilon'}{2})\bar{\mathbf{x}}|_I$ is $(1 - \frac{\varepsilon'}{2})^2$-feasible for $\boldsymbol{\mathcal{L}}^I$, and since $(1 - \frac{\varepsilon'}{2})^2 \ge (1 - \varepsilon')$, Part (a) follows.

Part (b). Let $\bar{\mathbf{x}}$ be an optimal solution for $\mathcal{L}$. Then again by Bernstein's inequality (Corollary 2.4)

$$\Pr\left(\sum_{t\in I} \boldsymbol{\pi}^t \bar{\mathbf{x}}^t \le \frac{s}{n}\mathsf{opt}(\mathcal{L}) - \tau\mathsf{opt}(\mathcal{L})\right) \le 2\exp\left(-\min\left\{\frac{\tau^2\mathsf{opt}(\mathcal{L})^2\vartheta}{8\rho(s/n)\mathsf{opt}(\mathcal{L})^2}, \frac{\tau\mathsf{opt}(\mathcal{L})\vartheta}{2\rho(s/n)\mathsf{opt}(\mathcal{L})}\right\}\right);$$

setting $\tau = (\varepsilon\rho/2)\sqrt{(s/n)}$ makes the right-hand side at most 2δ (this uses the fact $\rho \ge 1$).

Then using Part (a), with probability at least $1 - 4\delta$ we have $\bar{\mathbf{x}}|_I$ feasible for $\boldsymbol{\mathcal{L}}^I(1 - \varepsilon')$ and with value at least $(1 - \frac{\rho\varepsilon'}{2})\frac{s}{n}\mathsf{opt}(\boldsymbol{\mathcal{L}})$; this lower bounds the value of $\mathsf{opt}(\boldsymbol{\mathcal{L}}^I(1 - \varepsilon'))$ and Part (b) follows.

Now for the arguments about the upper bounds. Here we use the (ε, σ)-stability of $\mathcal{L}$. We consider the (random) dual $\boldsymbol{\mathcal{D}}^I$ of $\boldsymbol{\mathcal{L}}^I$ (recall from (B.19) the expression for the dual).

Part (d). Let $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\gamma}})$ be an optimal solution for $\boldsymbol{\mathcal{D}}(1 - \varepsilon)$. Notice that in every scenario this solution is feasible for $\boldsymbol{\mathcal{D}}^I(1 - \varepsilon')$, so taking expectations and using Strong Duality $\mathbb{E}[\boldsymbol{\mathcal{D}}^I(1 - \varepsilon')] \le \frac{s}{n}\mathsf{opt}(\mathcal{L}(1 - \varepsilon)) + (\varepsilon' - \varepsilon)\frac{s}{n}\langle\bar{\boldsymbol{\beta}}, \vartheta\mathbb{1}\rangle$. Using (ε, σ)-stability of $\boldsymbol{\mathcal{L}}$ we have $\mathsf{opt}(\mathcal{L}(1 - \varepsilon)) \le (1 + \sigma\varepsilon)\mathsf{opt}(\mathcal{L})$ and using Lemma B.1 $\langle\bar{\boldsymbol{\beta}}, \vartheta\mathbb{1}\rangle \le \sigma \cdot \mathsf{opt}(\mathcal{L})$. Since $\varepsilon' \ge \varepsilon$, this gives $\mathbb{E}[\boldsymbol{\mathcal{D}}^I(1 - \varepsilon')] \le \frac{s}{n}(1 + \sigma\varepsilon')\mathsf{opt}(\mathcal{L})$, and Strong Duality then gives Part (e) of the lemma.

Part (c). Again let $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\gamma}})$ be an optimal solution for $\boldsymbol{\mathcal{D}}(1 - \varepsilon)$. Notice that the above argument gives that in every scenario $\mathsf{opt}(\mathcal{L}^I(1-\varepsilon')) \le \frac{s}{n}\mathsf{opt}(\mathcal{L}(1-\varepsilon)) + (\varepsilon'-\varepsilon)\sigma\cdot\mathsf{opt}(\mathcal{L}) + (\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t - \mathbb{E}[\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t]) \le \frac{s}{n}(1+\sigma\varepsilon')\mathsf{opt}(\mathcal{L}) + (\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t - \mathbb{E}[\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t])$. Thus it suffices to show that with probability at least $1-\delta$, $\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t \le \mathbb{E}[\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t] + \varepsilon'\frac{s}{n}(\rho + \sigma)\mathsf{opt}(\mathcal{L})$.

For that, we need to control how large $\bar{\boldsymbol{\gamma}}$ can be. By the optimality of $\bar{\boldsymbol{\gamma}}$, notice that in each scenario and for all t we have $\bar{\boldsymbol{\gamma}}_t \le \max_j \{\boldsymbol{\pi}_j^t, \langle\bar{\boldsymbol{\beta}}, C_{\star,j}^t\rangle\}$, or else we could reduce $\bar{\boldsymbol{\gamma}}_t$ and obtain a feasible solution for $\boldsymbol{\mathcal{D}}(1-\varepsilon)$ with strictly smaller cost. Again from Lemma B.1 we have $\sum_i \bar{\boldsymbol{\beta}}_i \le \sigma \cdot \mathsf{opt}(\mathcal{L})/\vartheta$, and since the entries of $\mathbf{C}^t$ are at most 1, we get $\langle\bar{\boldsymbol{\beta}}, \mathbf{C}_{\star,j}^t\rangle \le \sigma \cdot \mathsf{opt}(\mathcal{L})/\vartheta$. Using our assumption that $\boldsymbol{\pi}_j^t \le \rho\mathsf{opt}(\mathcal{L})/\vartheta$, we get $\bar{\boldsymbol{\gamma}}_t \le (\sigma + \rho)\mathsf{opt}(\mathcal{L})/\vartheta$.

In addition, we claim that $\mathbb{E}[\sum_{t\in I}\bar{\boldsymbol{\gamma}}_t] \le \frac{s}{n}(1 + \sigma)\mathsf{opt}(\mathcal{L})$. Since $\mathsf{opt}(\mathcal{L}(1 - \varepsilon)) = \mathsf{opt}(\boldsymbol{\mathcal{L}}(1 - \varepsilon)) = \langle\bar{\boldsymbol{\alpha}}, \vartheta\mathbb{1}\rangle - (1 - \varepsilon)\langle\bar{\boldsymbol{\beta}}, \vartheta\mathbb{1}\rangle + \sum_t \bar{\boldsymbol{\gamma}}_t$, reorganizing we get $\sum_t \bar{\boldsymbol{\gamma}}_t \le \mathsf{opt}(\mathcal{L}(1 - \varepsilon)) + (1 - \varepsilon)\langle\bar{\boldsymbol{\beta}}, \vartheta\mathbb{1}\rangle$. Using

(ε,σ)-stability of $\mathcal{L}$ and again the bound from Lemma B.1, this gives $\sum_t \bar{\gamma}_t \le (1+\sigma)\mathsf{opt}(\mathcal{L})$, and the claim follows.

Then using Bernstein's inequality (Corollary 2.4), we get (letting $\mu = \mathbb{E}\big[\sum_{t\in I}\bar{\gamma}_t\big]$)

$$\Pr\left(\sum_{t\in I}\bar{\gamma}_t \ge \mu + \varepsilon\tau\mathsf{opt}(\boldsymbol{\mathcal{L}})\right) \le 2\exp\left(-\min\left\{\frac{\varepsilon^2\tau^2\mathsf{opt}(\mathcal{L})^2\vartheta}{8\mu(\rho+\sigma)\mathsf{opt}(\mathcal{L})}, \frac{\varepsilon\tau\mathsf{opt}(\mathcal{L})\vartheta}{2(\rho+\sigma)\mathsf{opt}(\mathcal{L})}\right\}\right)$$
$$\le 2\exp\left(-\frac{\log(1/\delta)}{(\rho+\sigma)}\min\left\{\frac{4\tau^2}{(s/n)(1+\sigma)}, \frac{16\tau}{\varepsilon}\right\}\right);$$

setting $\tau = (\rho+\sigma)\sqrt{(s/n)}$ makes the right-hand side at most 2δ (this uses the fact that $\rho \ge 1$ and $s \ge \varepsilon^2 n$). Since $\varepsilon'\frac{s}{n}(\rho+\sigma) = \varepsilon\tau$, this concludes the proof of the lemma.

D Proof of Theorem 6.1

First we verify that with good probability we have $\widehat{\mathsf{opt}} \in [\frac{1}{2}\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')), \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon'))]$ so we can employ Theorem 4.1 to obtain guarantees for the solution $\{\bar{\mathbf{x}}^t\}_{t=n/2+1}^n$ returned by **LPviaLB**.

Taking a union bound over Parts (b) and (c) of Lemma 5.5 and using the fact that $\sigma \ge 1$, we have that

$$\Pr\left[\left(1-\frac{\varepsilon'}{2}\right) \le \frac{\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\le\frac{n}{2}}(1-\varepsilon'))}{\frac{1}{2}\mathsf{opt}(\boldsymbol{\mathcal{L}})} \le (1+3\varepsilon'\sigma)\right] \ge 1-6\delta,$$

and the same holds for $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon'))$; let $\mathcal{E}$ be the event that these bounds hold for both, so we have $\Pr(\mathcal{E}) \ge 1-12\delta$. We can chain the bounds for $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{\le\frac{n}{2}}(1-\varepsilon'))$ and $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon'))$ to get that conditioned on $\mathcal{E}$

$$\frac{(1-\frac{\varepsilon'}{2})^2}{(1+3\varepsilon'\sigma)^2}\cdot\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')) \le \widehat{\mathsf{opt}} \le \mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')).$$

Now using the fact that $\epsilon \in (0,1)$ and $\sigma \ge 1$, it is easy to verify that we can lower bound the first factor as $\frac{(1-\frac{\varepsilon'}{2})^2}{(1+3\varepsilon'\sigma)^2} \ge 1-10\varepsilon'\sigma$. In addition, since $\sigma \le \frac{1}{20\varepsilon'}$, this in particular implies $\widehat{\mathsf{opt}} \ge \frac{1}{2}\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon'))$. Then under the event $\mathcal{E}$ the requirements of the Theorem 4.1 holds and thus taking a union bound over the event $\mathcal{E}$ and this theorem gives that with probability at least $1-13\delta$ the solution $\{\bar{\mathbf{x}}^t\}_{t=n/2+1}^n$ is $c_2\varepsilon$-feasible for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')$ (and thus is $(c_2\varepsilon+\varepsilon')$-feasible for $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}$) and has value

$$\sum_{t=n/2+1}^{n}\boldsymbol{\pi}^t\bar{\mathbf{x}}^t \ge (1-c_3\varepsilon)\,\widehat{\mathsf{opt}} \ge (1-c_3\varepsilon)(1-10\varepsilon'\sigma)\,\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon'))$$
$$\ge (1-c_5\varepsilon\sigma)\,\mathsf{opt}(\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')),$$

for some constant c_5, where the last inequality uses the fact $\sigma \ge 1$, which gives the desired guarantees. (Actually, in more detail notice that the event $\mathcal{E}$ only depends on what *set* of blocks (π^t, A^t, C^t) of the underlying LP $\mathcal{L}$ comprise the blocks in the random sampled LP $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')$ (this also determines the blocks that comprise the sampled LP $\boldsymbol{\mathcal{L}}^{\le\frac{n}{2}}(1-\varepsilon')$). Thus, we can first condition that the set of blocks in $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')$ is a specific one that satisfies $\mathcal{E}$, and since we do not fix their *order* they still appear in random order in the conditioned $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')$ and we can apply Theorem 4.1 to this conditioned $\boldsymbol{\mathcal{L}}^{>\frac{n}{2}}(1-\varepsilon')$; ranging over all such conditioning then gives the result.)

E Proof of Lemma 6.3

First we claim that with probability at least $1-6\delta$, the random LP $\boldsymbol{\mathcal{L}}^{S_i}(1-\varepsilon_i)$ is $(\varepsilon_i, 7\sigma)$-stable. Notice that we can employ Lemma 5.5(c) to $\mathcal{L}$ with parameter ε' set to $2\varepsilon_i$ (in particular, $2\varepsilon_i\sqrt{|S_i|/n} \geq \varepsilon$ plus Lemma B.1 implies that $\mathcal{L}$ is $(2\varepsilon_i\sqrt{|S_i|/n}, \sigma)$-stable). Thus, Lemma 5.5(c) guarantees that with probability at least $1-2\delta$, $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i}(1-2\varepsilon_i)) \leq (1+6\sigma\varepsilon_i)\frac{|S_i|}{n}\mathsf{opt}(\mathcal{L})$. Moreover, applying Lemma 5.5(b) to $\mathcal{L}$ with $\varepsilon' = \varepsilon_i$ we get that with probability at least $1-4\delta$, $\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i}(1-\varepsilon_i)) \geq (1-\frac{\varepsilon_i}{2})\frac{|S_i|}{n}\mathsf{opt}(\mathcal{L})$. Putting these inequalities together, we get that with probability at least $1-6\delta$

$$\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i}(1-2\varepsilon_i)) \leq \frac{(1+6\sigma\varepsilon_i)}{(1-\frac{\varepsilon_i}{2})}\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i}(1-\varepsilon_i)) \leq (1+7\sigma\varepsilon_i)\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i}(1-\varepsilon_i)),$$

where the last inequality uses $\sigma \geq 1$ and that $\varepsilon_i \leq \sqrt{\varepsilon}$ is at most a sufficiently small constant (for example $\sqrt{\varepsilon} \leq \frac{1}{7}$ suffices). This proves the claim.

Second, notice that whenever the claim above holds (i.e. $\boldsymbol{\mathcal{L}}^{S_i}(1-\varepsilon_i)$ is $(\varepsilon_i, 7\sigma)$-stable) we can apply Theorem 6.1 with parameters $(\varepsilon_i, \delta, 7\sigma)$ to obtain guarantees for the solution $(0, \ldots, 0, \bar{\mathbf{x}}^{|S_{i-1}|+1}, \ldots, \bar{\mathbf{x}}^{|S_i|})$. So employing the claim above, Theorem 6.1 and union bound, we get that with probability at least $1-19\delta$ this solution is $(c_4\varepsilon_i)$-feasible for $\boldsymbol{\mathcal{L}}^{S_i\setminus S_{i-1}}(1-\varepsilon_i)$ (and thus $((c_4+1)\varepsilon_i)$-feasible for $\boldsymbol{\mathcal{L}}^{S_i\setminus S_{i-1}}$) and has value $\sum_{t=|S_{i-1}|+1}^{|S_i|}\boldsymbol{\pi}\bar{\mathbf{x}}^t \geq (1-c_5 7\sigma\varepsilon_i)\mathsf{opt}(\boldsymbol{\mathcal{L}}^{S_i\setminus S_{i-1}}(1-\sqrt{2}\varepsilon_i))$. This concludes the proof of the lemma. (As noted at the end of Theorem 6.1, formally we should condition on the *set* of blocks of $\boldsymbol{\mathcal{L}}^{S_i}$ being such that the claim above holds and then apply Theorem 6.1 on the conditioned LP (which still has its blocks coming in random order), and then take an average over all such conditionings.)

F Proof of Lemma 7.7

The following will be a more convenient way of seeing the random order process: the random restricted LP $\boldsymbol{\mathcal{L}}^I$ (i.e. columns are in random order and we are focusing on the ones indexed by I) has the same distribution as the LP $\mathcal{L}^{\mathbf{I}}$ for $\mathbf{I}$ a uniformly chosen $|I|$-subset of $[n]$ (i.e. sampling $|I|$-columns from $\mathcal{L}$ without replacement).

The following lemma is the main part of the argument.

Lemma F.1. *Let $\varepsilon \leq 1/20$. Given $\mathcal{L} \in \mathsf{PLP}_n$ with width at least $40B$, let τ be an $\mathsf{opt}(\mathcal{L})$-good skimming threshold for $\mathcal{L}$. Let $\mathbf{K}$ be a random pn-subset of $[n]$ for $p \in [\frac{1}{4}, \frac{3}{4}]$. Then with probability at least $1-2\delta$,*

$$\mathsf{rhs}(\mathit{skim}_\tau(\mathcal{L}^{\mathbf{K}})) \in (1\pm\varepsilon)\, p\, \mathsf{rhs}(\mathit{skim}_\tau(\mathcal{L}))$$

and with probability at least $1-8\delta$,

$$\mathsf{opt}(\mathit{skim}_\tau(\mathcal{L}^{\mathbf{K}})) \in (1\pm\varepsilon)p\,\mathsf{opt}(\mathit{skim}_\tau(\mathcal{L})) \pm \varepsilon 64 pB\tau.$$

Proof. Let $H = \{t : \pi_t > \tau\}$ be the set of high-value items. If A^t is the t^{th} column of A, we get that

$$\begin{aligned} b_L &:= \mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L})) \;=\; \mathsf{rhs}(\mathcal{L}) - \textstyle\sum_{t\in H} A^t \;; \\ b_L^{\mathbf{K}} &:= \mathsf{rhs}(\mathsf{skim}_\tau(\mathcal{L}^{\mathbf{K}})) \;=\; \mathsf{rhs}(\mathcal{L}^{\mathbf{K}}) - \textstyle\sum_{t\in\mathbf{K}:\pi_t>\tau} A^t \;=\; p\,\mathsf{rhs}(\mathcal{L}) - \textstyle\sum_{t\in H\cap\mathbf{K}} A^t. \end{aligned}$$

Since $\mathbf{K}$ is a random pn-subset of the columns, the expectation of the second line is precisely p times the first line. To obtain concentration, fix a coordinate i. The second property of a good skimming threshold implies the high-value items occupy at most half of the budget: $\sum_{t\in H} A_i^t \leq \frac{\mathsf{rhs}(\mathcal{L})_i}{2} \leq (b_L)_i$.

This gives $\mathbb{E}\big[\sum_{t\in H\cap\mathbf{K}} A_i^t\big] = p\sum_{t\in H} A_i^t \le p\,(b_L)_i$ and Bernstein's inequality Corollary 2.4 gives that $\sum_{t\in H\cap\mathbf{K}} A_i^t = p\sum_{t\in H} A_i^t \pm \varepsilon\, p\,(b_L)_i$ with probability at least

$$1 - 2\exp\left\{-\frac{(\varepsilon\, p\,(b_L)_i)^2}{(4+\varepsilon)p\,(b_L)_i}\right\} \ge 1 - 2\exp\left\{-\frac{\varepsilon^2 p\,(b_L)_i}{5}\right\} \ge 1 - \frac{2\delta}{m},$$

where the last inequality uses the width lower bound $(b_L)_i \ge \frac{1}{2}\,\mathsf{rhs}(\mathcal{L})_i \ge 20B$ (recall we have used the normalization $A_i^t \in [0,1]$). A union bound over all coordinates i then gives that $b_L^{\mathbf{K}} \in p\,b_L(1\pm\varepsilon)$ with probability at least $1-2\delta$, giving the first part of the lemma.

For the estimate of the optimal value, let $L := [n]\setminus H$ be the low-value items. Let $\mathsf{skim}_\tau(\mathcal{L})\big(\frac{1}{p}\cdot b_L^{\mathbf{K}}\big)$ be the LP $\mathsf{skim}_\tau(\mathcal{L})$ with the right-hand side replaced by $\frac{1}{p}\cdot b_L^{\mathbf{K}}$ and notice that

$$\mathsf{skim}_\tau(\mathcal{L}^{\mathbf{K}}) = \max\big\{\textstyle\sum_{t\in L\cap\mathbf{K}} \pi_t x_t \mid \sum_{t\in\mathbf{K}} A^t x_t \le b_L^{\mathbf{K}}\big\} = \big(\mathsf{skim}_\tau(\mathcal{L})\big(\tfrac{1}{p}\cdot b_L^{\mathbf{K}}\big)\big)^{\mathbf{K}}. \tag{F.20}$$

The idea is to apply Lemma 5.5 to bound $\mathsf{skim}_\tau(\mathcal{L})\big(\frac{1}{p}\cdot b_L^{\mathbf{K}}\big)$, but unfortunately the right-hand side of this LP is random; but since from the previous part $\frac{1}{p}\cdot b_L^{\mathbf{K}} \in b_L(1\pm\varepsilon)$ with good probability, we replace the right-hand side with the upper bound $(1+\varepsilon)b_L$ (resp. lower bound $(1-\varepsilon)b_L$) to upper (resp. lower) bound the optimal value. So define $\mathcal{L}_+ = \mathsf{skim}_\tau(\mathcal{L})\big(b_L(1+\varepsilon)\big)$ and employ Lemma 5.5(c) to $\mathcal{L}_+$ with $\rho = \frac{32\tau B}{\mathsf{opt}(\mathcal{L}_+)}$ to get that with probability at least $1-2\delta$

$$\mathsf{opt}(\mathcal{L}_+^{\mathbf{K}}) \le \left(1 + \frac{\varepsilon 64\tau B}{\mathsf{opt}(\mathcal{L}_+)}\right) p\,\mathsf{opt}(\mathcal{L}_+) \le (1+\varepsilon)p\,\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L})) + \varepsilon 64 pB\tau, \tag{F.21}$$

where again $\mathsf{opt}(\mathcal{L}_+) \le (1+\varepsilon)\,\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}))$ since an optimal solution to $\mathcal{L}_+$ can be scaled by $\frac{1}{1+\varepsilon}$ to give a feasible solution to $\mathsf{skim}_\tau(\mathcal{L})$. Moreover, from the first part of this lemma, with probability at least $1-2\delta$ we have $\frac{1}{p}\cdot b_L^{\mathbf{K}} \le b_L(1+\varepsilon)$, in which case $\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L})(\frac{1}{p}\cdot b_L^{\mathbf{K}})^{\mathbf{K}}) \le \mathsf{opt}(\mathcal{L}_+^{\mathbf{K}})$. Thus, together with inequalities (F.20) and (F.21), a union bound gives that with probability at least $1-4\delta$

$$\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}^{\mathbf{K}})) \le (1+\varepsilon)p\,\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L})) + \varepsilon 64 pB\tau.$$

Using the same reasoning, the lower bound

$$\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L}^{\mathbf{K}})) \ge (1-\varepsilon)p\,\mathsf{opt}(\mathsf{skim}_\tau(\mathcal{L})) - \varepsilon 64 pB\tau$$

holds with probability at least $1-4\delta$. Taking a union bound over these upper and lower bounds then concludes the proof of the lemma. □

Proof of Lemma 7.7. Let $\mathbf{I}$ be a random $\frac{n}{3}$-subset of $[n]$ and let $\mathbf{J} = [n]\setminus\mathbf{I}$, and recall $\mathsf{skim}_\tau(\boldsymbol{\mathcal{L}}^{\le\frac{n}{3}})$ has the same distribution as $\mathsf{skim}_\tau(\mathcal{L}^{\mathbf{I}})$ and $\mathsf{skim}_\tau(\boldsymbol{\mathcal{L}}^{>\frac{n}{3}})$ has the same distribution as $\mathsf{skim}_\tau(\mathcal{L}^{\mathbf{J}})$. Thus, employing Lemma F.1 with $\mathbf{K}=\mathbf{I}$ and then $\mathbf{K}=\mathbf{J}$ and employing a union bound, we get that with probability at least $1-4\delta$

$$\mathsf{rhs}\big(\mathsf{skim}_\tau(\boldsymbol{\mathcal{L}}^{\le\frac{n}{3}})\big) \in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right]\frac{1}{2}\,\mathsf{rhs}\big(\mathsf{skim}_\tau(\boldsymbol{\mathcal{L}}^{>\frac{n}{3}})\big) \subseteq (1\pm 3\varepsilon)\,\frac{1}{2}\,\mathsf{rhs}\big(\mathsf{skim}_\tau(\boldsymbol{\mathcal{L}}^{>\frac{n}{3}})\big),$$

where the last inequality uses $\varepsilon \le \frac{1}{3}$.

The bounds for $\mathsf{opt}\big(\mathsf{skim}_\tau(\boldsymbol{\mathcal{L}}^{\le\frac{n}{3}})\big)$ follow similarly. This concludes the proof. □

G Proof of Theorem 7.11

We prove only the first part of the theorem since the second directly follows from it. Let $\mathcal{E}$ be the event that the the guarantee of Theorem 7.10 holds for all $i = 1, \dots, \log \frac{1}{\varepsilon}$. Since each such guarantee holds with probability at least $1 - 15\delta$, a union bound gives that $\mathcal{E}$ holds with probability at least $1 - 15\delta \log \frac{1}{\varepsilon}$. As in the proof of Theorem 6.2, it is easy to see that under $\mathcal{E}$ the returned solution $\{\bar{\mathbf{x}}_t\}_t$ is feasible for $\boldsymbol{\mathcal{L}}$.

Now we analyze the value val of the solution. To simplify the notation, let $\mathsf{opt}(I) = \mathsf{opt}(\boldsymbol{\mathcal{L}}^I)$ for any subset $I \subseteq [n]$. Again, conditioned on $\mathcal{E}$ we can put together the guarantees from Theorem 7.10 for each phase of **mDLA** to get (recall in phase i we use the approximation parameter $\varepsilon_i = \varepsilon\sqrt{n/|S_i|}$)

$$\mathbb{E}[\mathsf{val} \,|\, \mathcal{E}] \geq \sum_{i=1}^{\log \varepsilon^{-1}} \mathbb{E}\big[\mathsf{opt}(S_i \setminus S_{i-1}) - c_6 \varepsilon_i \mathsf{opt}(S_i) \,|\, \mathcal{E}\big]. \tag{G.22}$$

First we analyze the unconditioned value of the right-hand side of this expression.

Lemma G.1. $\sum_{i=1}^{\log \varepsilon^{-1}} \mathbb{E}[\mathsf{opt}(S_i \setminus S_{i-1})] - c_6 \sum_{i=1}^{\log \varepsilon^{-1}} \varepsilon_i \, \mathbb{E}[\mathsf{opt}(S_i)] \geq (1 - 2\varepsilon - 2\delta \log \varepsilon^{-1} - 4c_6\varepsilon)\, \mathsf{opt}(\mathcal{L})$.

Proof. The last term in the left-hand side is at most $4c_6\varepsilon\, \mathsf{opt}(\mathcal{L})$: Lemma 5.5(d) gives $\mathbb{E}[\mathsf{opt}(S_i)] \leq \frac{|S_i|}{n}\mathsf{opt}(\mathcal{L})$ (since our LP without covering constraints is trivially $(\varepsilon', 0)$-stable), so we get

$$c_6 \sum_{i=1}^{\log \varepsilon^{-1}} \varepsilon_i \, \mathbb{E}[\mathsf{opt}(S_i)] \leq c_6 \sum_{i=1}^{\log \varepsilon^{-1}} \varepsilon_i \, \frac{|S_i|}{n} \, \mathsf{opt}(\mathcal{L}) = c_6 \sum_{i=1}^{\log \varepsilon^{-1}} \varepsilon_i \, \varepsilon \, 2^i \, \mathsf{opt}(\mathcal{L}) \leq 4c_6 \varepsilon \, \mathsf{opt}(\mathcal{L}),$$

where the last inequality follows from the bound $\sum_{i=1}^{\log \varepsilon^{-1}} \varepsilon_i 2^{i-1} \leq 2$ from equation (6.14).

To show that the first term in the left-hand side is approximately $\mathsf{opt}(\mathcal{L})$ is a little more complicated, since without good bounds on the ratios $\frac{\pi_t}{\mathsf{opt}(\mathcal{L})}$ we do not have bounds showing that $\mathsf{opt}(S_i \setminus S_{i-1}) \approx \frac{|S_i|\setminus|S_{i-1}|}{n}\, \mathsf{opt}(\mathcal{L})$ even in expectation; we need to look at all the terms $\mathsf{opt}(S_i \setminus S_{i-1})$ simultaneously.

Let $\mathbf{x}^*$ be a optimal solution for $\boldsymbol{\mathcal{L}}$ and let $\mathbf{x}^*(i)$ be the solution scaled down by $(1 - \frac{\varepsilon_i}{2})$ and restricted to the $2^{i-1}\varepsilon n$ columns in $S_i \setminus S_{i-1}$. From Lemma 5.5(a) we have that $\mathbf{x}^*(i)$ is feasible for $\boldsymbol{\mathcal{L}}^{S_i \setminus S_{i-1}}$ with probability at least $1 - 2\delta$. Let $\mathcal{G}$ be the good event that for all i we have $\mathbf{x}^*(i)$ feasible for $\boldsymbol{\mathcal{L}}^{S_i \setminus S_{i-1}}$, which holds with probability at least $1 - 2(\log \varepsilon^{-1})\delta$ by a union bound. Then for every scenario in $\mathcal{G}$ we have $\mathsf{opt}(S_i \setminus S_{i-1}) \geq (1 - \frac{\varepsilon_i}{2}) \sum_{t \in S_i \setminus S_{i-1}} \boldsymbol{\pi}_t \mathbf{x}^*_t$, so adding over all $i \geq 1$ gives

$$\sum_{i \geq 1} \mathsf{opt}(S_i \setminus S_{i-1}) \geq \mathsf{opt}(\mathcal{L}) - \sum_{i \geq 1} \frac{\varepsilon_i}{2} \sum_{t \in S_i \setminus S_{i-1}} \boldsymbol{\pi}_t \mathbf{x}^*_t - \sum_{t \in S_0} \boldsymbol{\pi}_t \mathbf{x}^*_t.$$

Then letting $\mathbb{1}_{\mathcal{G}}$ be the indicator of the event $\mathcal{G}$ and using the non-negativity of the objective function, we get that

$$\mathbb{E}\left[\sum_{i \geq 1} \mathsf{opt}(S_i \setminus S_{i-1})\right] \geq \mathbb{E}\left[\mathbb{1}_{\mathcal{G}} \cdot \sum_{i \geq 1} \mathsf{opt}(S_i \setminus S_{i-1})\right] \geq \mathbb{E}[\mathbb{1}_{\mathcal{G}} \cdot \mathsf{opt}(\mathcal{L})] - \sum_{i \geq 1} \frac{\varepsilon_i}{2} \sum_{t \in S_i \setminus S_{i-1}} \mathbb{E}[\boldsymbol{\pi}_t \mathbf{x}^*_t] - \sum_{t \in S_0} \mathbb{E}[\boldsymbol{\pi}_t \mathbf{x}^*_t].$$

Clearly $\mathbb{E}[\mathbb{1}_{\mathcal{G}} \cdot \mathsf{opt}(\mathcal{L})] = \Pr(\mathcal{G}) \cdot \mathsf{opt}(\mathcal{L}) \geq (1 - 2\delta \log \varepsilon^{-1})\, \mathsf{opt}(\mathcal{L})$. Moreover, by the random order model we have that for all t, $\mathbb{E}[\boldsymbol{\pi}_t \mathbf{x}^*_t] = \frac{\mathsf{opt}(\mathcal{L})}{n}$. Thus, the above bound becomes

$$\mathbb{E}\left[\sum_{i \geq 1} \mathsf{opt}(S_i \setminus S_{i-1})\right] \geq (1 - 2\delta \log \varepsilon^{-1})\, \mathsf{opt}(\mathcal{L}) - \frac{\mathsf{opt}(\mathcal{L})}{n} \sum_{i \geq 1} \frac{\varepsilon_i \, \varepsilon \, 2^{i-1} n}{2} - \varepsilon \mathsf{opt}(\mathcal{L}) \geq (1 - 2\varepsilon - 2\delta \log \varepsilon^{-1})\, \mathsf{opt}(\mathcal{L}),$$

where again we have used equation (6.14) in the last step. This concludes the proof of the lemma. □

To handle the *conditioned* right-hand side of (G.22) we observe the following: with probability 1 we have $\mathsf{opt}(S_0) + \sum_{i=1}^{\log \varepsilon^{-1}} \mathsf{opt}(S_i \setminus S_{i-1}) \leq \mathsf{opt}(\mathcal{L})$, since concatenating the optimal solutions for $\boldsymbol{\mathcal{L}}^{S_0}$, $\boldsymbol{\mathcal{L}}^{S_1 \setminus S_0}$, ..., gives a feasible solution for $\boldsymbol{\mathcal{L}}$. For convenience, let $\mathsf{val}_i := \mathsf{opt}(S_i \setminus S_{i-1}) - c_6 \varepsilon_i \mathsf{opt}(S_i)$. Then using the decomposition $\mathbb{E}[\sum_{i\geq 1} \mathsf{val}_i] = \mathbb{E}[\sum_{i\geq 1} \mathsf{val}_i \,|\, \mathcal{E}] \Pr(\mathcal{E}) + \mathbb{E}[\sum_{i\geq 1} \mathsf{val}_i \,|\, \bar{\mathcal{E}}] \Pr(\bar{\mathcal{E}})$, inequality (G.22) gives

$$\begin{aligned}
\mathbb{E}[\mathsf{val} \,|\, \mathcal{E}] &\geq \sum_{i\geq 1} \mathbb{E}[\mathsf{val}_i \,|\, \mathcal{E}] \geq \sum_{i\geq 1} \mathbb{E}[\mathsf{val}_i] - \Pr(\bar{\mathcal{E}}) \sum_{i\geq 1} \mathbb{E}[\mathsf{val}_i \,|\, \bar{\mathcal{E}}] \\
&\geq \sum_{i\geq 1} \mathbb{E}[\mathsf{val}_i] - \Pr(\bar{\mathcal{E}}) \left[\mathbb{E}[\mathsf{opt}(S_0) \,|\, \bar{\mathcal{E}}] + \sum_{i\geq 1} \mathbb{E}[\mathsf{opt}(S_i \setminus S_{i-1}) \,|\, \bar{\mathcal{E}}] \right] \\
&\geq \sum_{i\geq 1} \mathbb{E}[\mathsf{val}_i] - \Pr(\bar{\mathcal{E}})\, \mathsf{opt}(\mathcal{L}).
\end{aligned}$$

Lower bounding the first term using Lemma G.1 and recalling that $\Pr(\bar{\mathcal{E}}) \leq 15\delta \log \varepsilon^{-1}$, gives $\mathbb{E}[\mathsf{val} \,|\, \mathcal{E}] \geq (1 - 2\varepsilon - 17\delta \log \varepsilon^{-1} - 4c_6\varepsilon)\, \mathsf{opt}(\mathcal{L})$. This concludes the proof of Theorem 7.11.